Exhibit 2

Cisco	Arista
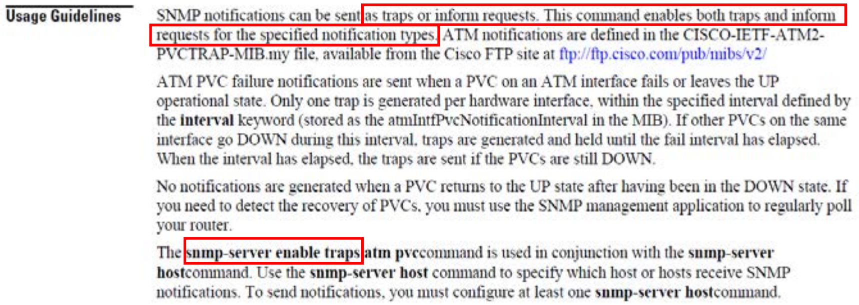 **Usage Guidelines** SNMP notifications can be sent as traps or inform requests. This command enables both traps and inform requests for the specified notification types. ATM notifications are defined in the CISCO-IETF-ATM2-PVCTRAP-MIB.my file, available from the Cisco FTP site at ftp://ftp.cisco.com/pub/mibs/v2/ ATM PVC failure notifications are sent when a PVC on an ATM interface fails or leaves the UP operational state. Only one trap is generated per hardware interface, within the specified interval defined by the **interval** keyword (stored as the atmIntfPvcNotificationInterval in the MIB). If other PVCs on the same interface go DOWN during this interval, traps are generated and held until the fail interval has elapsed. When the interval has elapsed, the traps are sent if the PVCs are still DOWN. No notifications are generated when a PVC returns to the UP state after having been in the DOWN state. If you need to detect the recovery of PVCs, you must use the SNMP management application to regularly poll your router. The **snmp-server enable traps atm pvc**command is used in conjunction with the **snmp-server host**command. Use the **snmp-server host** command to specify which host or hosts receive SNMP notifications. To send notifications, you must configure at least one **snmp-server host**command. Cisco IOS Asynchronous Transfer Mode Command Reference (2011), at 535	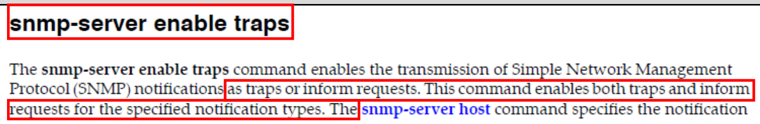 **snmp-server enable traps** The **snmp-server enable traps** command enables the transmission of Simple Network Management Protocol (SNMP) notifications as traps or inform requests. This command enables both traps and inform requests for the specified notification types. The **snmp-server host** command specifies the notification Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1918

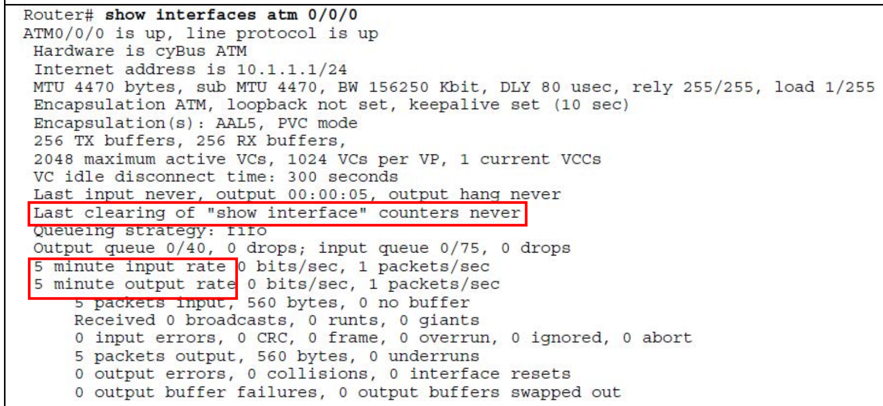
`Router# show interfaces atm 0/0/0`
`ATM0/0/0 is up, line protocol is up`
` Hardware is cyBus ATM`
` Internet address is 10.1.1.1/24`
` MTU 4470 bytes, sub MTU 4470, BW 156250 Kbit, DLY 80 usec, rely 255/255, load 1/255`
` Encapsulation ATM, loopback not set, keepalive set (10 sec)`
` Encapsulation(s): AAL5, PVC mode`
` 256 TX buffers, 256 RX buffers,`
` 2048 maximum active VCs, 1024 VCs per VP, 1 current VCCs`
` VC idle disconnect time: 300 seconds`
` Last input never, output 00:00:05, output hang never`
` Last clearing of "show interface" counters never`
` Queueing strategy: fifo`
` Output queue 0/40, 0 drops; input queue 0/75, 0 drops`
` 5 minute input rate 0 bits/sec, 1 packets/sec`
` 5 minute output rate 0 bits/sec, 1 packets/sec`
`     5 packets input, 560 bytes, 0 no buffer`
`     Received 0 broadcasts, 0 runts, 0 giants`
`     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort`
`     5 packets output, 560 bytes, 0 underruns`
`     0 output errors, 0 collisions, 0 interface resets`
`     0 output buffer failures, 0 output buffers swapped out`
 	
Cisco IOS Asynchronous Transfer Mode Command Reference (2011), at 476
	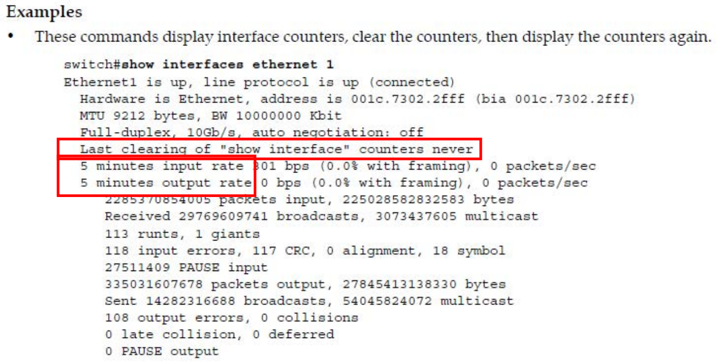
	Examples
	- These commands display interface counters, clear the counters, then display the counters again.
	`switch#show interfaces ethernet 1`
	`Ethernet1 is up, line protocol is up (connected)`
	`  Hardware is Ethernet, address is 001c.7302.2fff (bia 001c.7302.2fff)`
	`  MTU 9212 bytes, BW 10000000 Kbit`
	`  Full-duplex, 10Gb/s, auto negotiation: off`
	`  Last clearing of "show interface" counters never`
	`  5 minutes input rate 301 bps (0.0% with framing), 0 packets/sec`
	`  5 minutes output rate 0 bps (0.0% with framing), 0 packets/sec`
	`     2285370854005 packets input, 225028582832583 bytes`
	`     Received 29769609741 broadcasts, 3073437605 multicast`
	`     113 runts, 1 giants`
	`     118 input errors, 117 CRC, 0 alignment, 18 symbol`
	`     27511409 PAUSE input`
	`     335031607678 packets output, 27845413138330 bytes`
	`     Sent 14282316688 broadcasts, 54045824072 multicast`
	`     108 output errors, 0 collisions`
	`     0 late collision, 0 deferred`
	`     0 PAUSE output`
	
	Arista Configuration Guide v. 4.13.6F (4/14/2014), at 637

Cisco	Arista
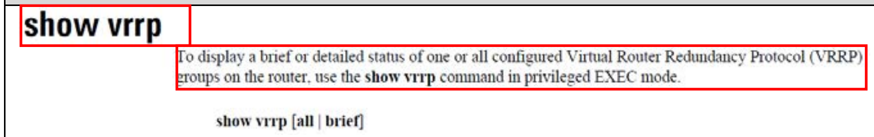 **show vrrp** To display a brief or detailed status of one or all configured Virtual Router Redundancy Protocol (VRRP) groups on the router, use the **show vrrp** command in privileged EXEC mode. **show vrrp [all | brief]** Cisco IOS IP Application Services Command Reference (2011), at 71	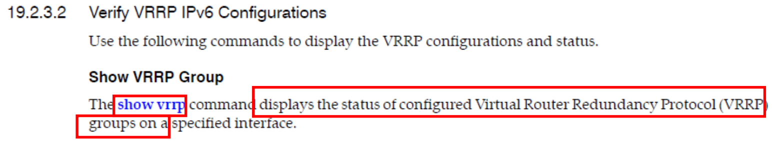 19.2.3.2 Verify VRRP IPv6 Configurations Use the following commands to display the VRRP configurations and status. **Show VRRP Group** The show vrrp command displays the status of configured Virtual Router Redundancy Protocol (VRRP) groups on a specified interface. Arista Configuration Guide v. 4.13.6F (4/14/2014), at 879
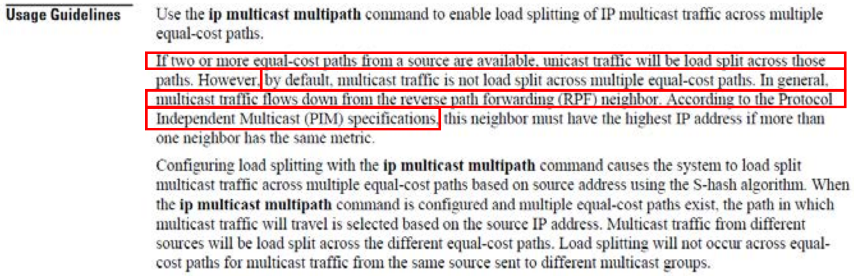 **Usage Guidelines** Use the **ip multicast multipath** command to enable load splitting of IP multicast traffic across multiple equal-cost paths. If two or more equal-cost paths from a source are available, unicast traffic will be load split across those paths. However, by default, multicast traffic is not load split across multiple equal-cost paths. In general, multicast traffic flows down from the reverse path forwarding (RPF) neighbor. According to the Protocol Independent Multicast (PIM) specifications, this neighbor must have the highest IP address if more than one neighbor has the same metric. Configuring load splitting with the **ip multicast multipath** command causes the system to load split multicast traffic across multiple equal-cost paths based on source address using the S-hash algorithm. When the **ip multicast multipath** command is configured and multiple equal-cost paths exist, the path in which multicast traffic will travel is selected based on the source IP address. Multicast traffic from different sources will be load split across the different equal-cost paths. Load splitting will not occur across equal-cost paths for multicast traffic from the same source sent to different multicast groups. Cisco IOS IP Multicast Command Reference (2011), at 293	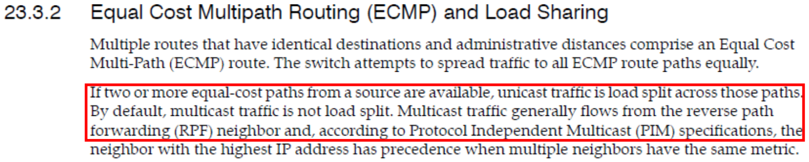 23.3.2 Equal Cost Multipath Routing (ECMP) and Load Sharing Multiple routes that have identical destinations and administrative distances comprise an Equal Cost Multi-Path (ECMP) route. The switch attempts to spread traffic to all ECMP route paths equally. If two or more equal-cost paths from a source are available, unicast traffic is load split across those paths. By default, multicast traffic is not load split. Multicast traffic generally flows from the reverse path forwarding (RPF) neighbor and, according to Protocol Independent Multicast (PIM) specifications, the neighbor with the highest IP address has precedence when multiple neighbors have the same metric. Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1191
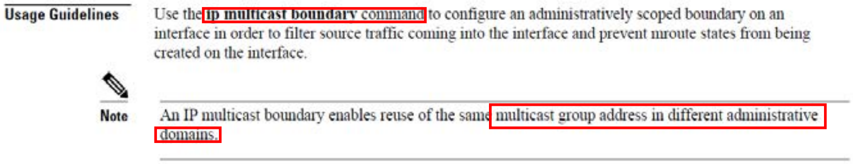 **Usage Guidelines** Use the **ip multicast boundary** command to configure an administratively scoped boundary on an interface in order to filter source traffic coming into the interface and prevent mroute states from being created on the interface. **Note** An IP multicast boundary enables reuse of the same multicast group address in different administrative domains. Cisco IOS IP Multicast Command Reference (2011), at 264	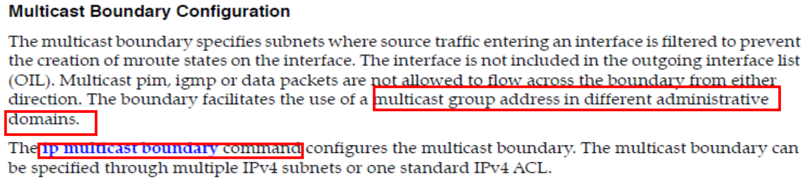 **Multicast Boundary Configuration** The multicast boundary specifies subnets where source traffic entering an interface is filtered to prevent the creation of mroute states on the interface. The interface is not included in the outgoing interface list (OIL). Multicast pim, igmp or data packets are not allowed to flow across the boundary from either direction. The boundary facilitates the use of a multicast group address in different administrative domains. The ip multicast boundary command configures the multicast boundary. The multicast boundary can be specified through multiple IPv4 subnets or one standard IPv4 ACL. Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1704

Cisco	Arista
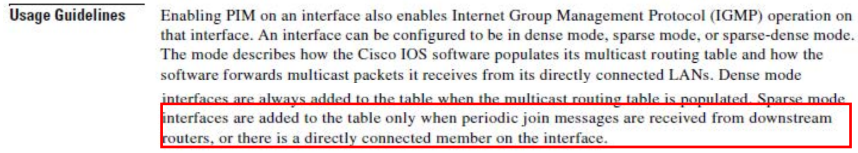 **Usage Guidelines** Enabling PIM on an interface also enables Internet Group Management Protocol (IGMP) operation on that interface. An interface can be configured to be in dense mode, sparse mode, or sparse-dense mode. The mode describes how the Cisco IOS software populates its multicast routing table and how the software forwards multicast packets it receives from its directly connected LANs. Dense mode interfaces are always added to the table when the multicast routing table is populated. Sparse mode interfaces are added to the table only when periodic join messages are received from downstream routers, or there is a directly connected member on the interface. Cisco IOS IP Multicast Command Reference (2008), at IMC-233–34	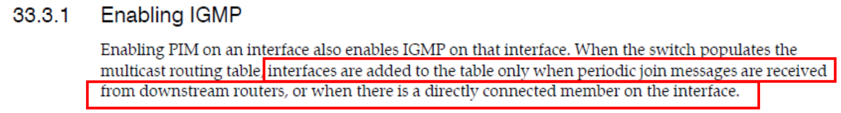 33.3.1 Enabling IGMP Enabling PIM on an interface also enables IGMP on that interface. When the switch populates the multicast routing table, interfaces are added to the table only when periodic join messages are received from downstream routers, or when there is a directly connected member on the interface. Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1726
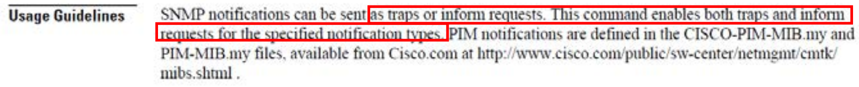 **Usage Guidelines** SNMP notifications can be sent as traps or inform requests. This command enables both traps and inform requests for the specified notification types. PIM notifications are defined in the CISCO-PIM-MIB.my and PIM-MIB.my files, available from Cisco.com at http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml . Cisco IOS IP Multicast Command Reference (2011), at 742	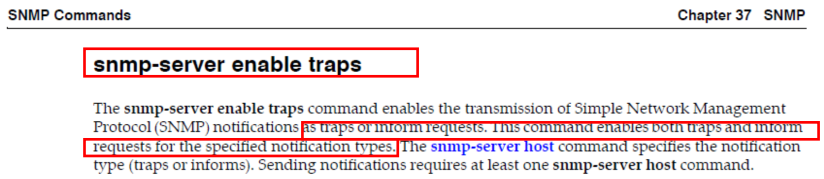 SNMP Commands Chapter 37 SNMP **snmp-server enable traps** The **snmp-server enable traps** command enables the transmission of Simple Network Management Protocol (SNMP) notifications as traps or inform requests. This command enables both traps and inform requests for the specified notification types. The snmp-server host command specifies the notification type (traps or informs). Sending notifications requires at least one **snmp-server host** command. Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1918
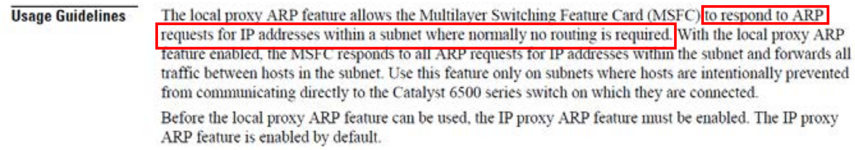 **Usage Guidelines** The local proxy ARP feature allows the Multilayer Switching Feature Card (MSFC) to respond to ARP requests for IP addresses within a subnet where normally no routing is required. With the local proxy ARP feature enabled, the MSFC responds to all ARP requests for IP addresses within the subnet and forwards all traffic between hosts in the subnet. Use this feature only on subnets where hosts are intentionally prevented from communicating directly to the Catalyst 6500 series switch on which they are connected. Before the local proxy ARP feature can be used, the IP proxy ARP feature must be enabled. The IP proxy ARP feature is enabled by default. Cisco IOS IP Addressing Services Command Reference (2011), at 394	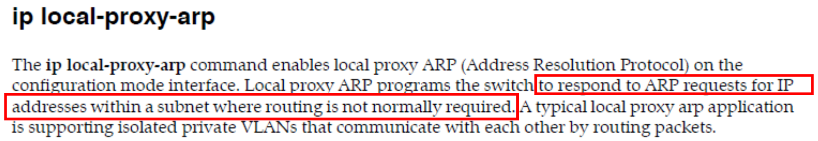 **ip local-proxy-arp** The **ip local-proxy-arp** command enables local proxy ARP (Address Resolution Protocol) on the configuration mode interface. Local proxy ARP programs the switch to respond to ARP requests for IP addresses within a subnet where routing is not normally required. A typical local proxy arp application is supporting isolated private VLANs that communicate with each other by routing packets. Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1231
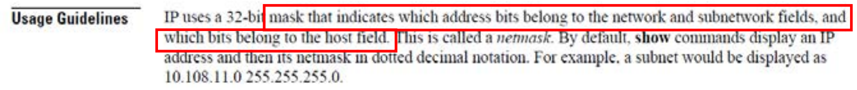 **Usage Guidelines** IP uses a 32-bit mask that indicates which address bits belong to the network and subnetwork fields, and which bits belong to the host field. This is called a *netmask*. By default, **show** commands display an IP address and then its netmask in dotted decimal notation. For example, a subnet would be displayed as 10.108.11.0 255.255.255.0. Cisco IOS IP Addressing Services Command Reference (2011), at 452	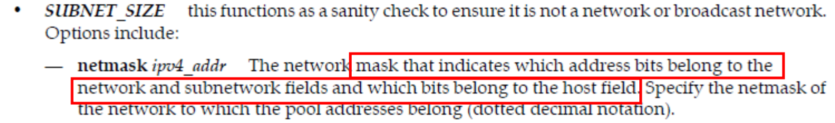 • *SUBNET_SIZE* this functions as a sanity check to ensure it is not a network or broadcast network. Options include: — **netmask** *ipv4_addr* The network mask that indicates which address bits belong to the network and subnetwork fields and which bits belong to the host field. Specify the netmask of the network to which the pool addresses belong (dotted decimal notation). Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1233

Cisco	Arista
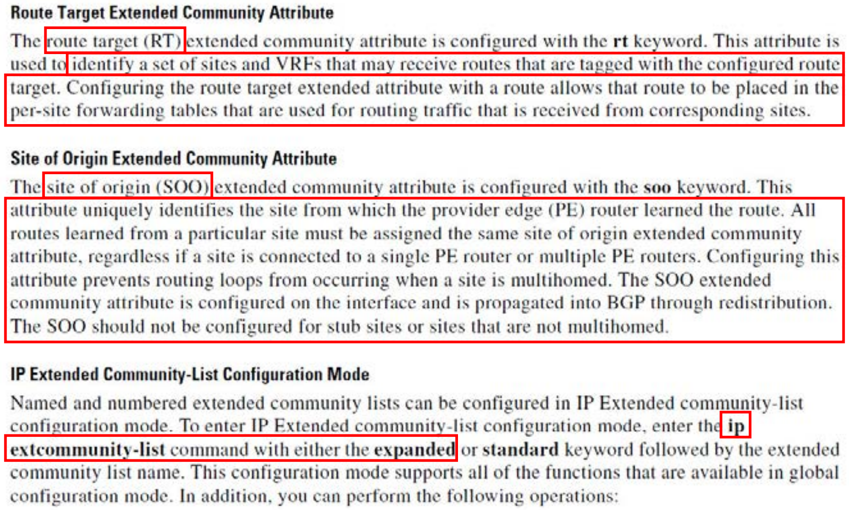 **Route Target Extended Community Attribute** The route target (RT) extended community attribute is configured with the **rt** keyword. This attribute is used to identify a set of sites and VRFs that may receive routes that are tagged with the configured route target. Configuring the route target extended attribute with a route allows that route to be placed in the per-site forwarding tables that are used for routing traffic that is received from corresponding sites. **Site of Origin Extended Community Attribute** The site of origin (SOO) extended community attribute is configured with the **soo** keyword. This attribute uniquely identifies the site from which the provider edge (PE) router learned the route. All routes learned from a particular site must be assigned the same site of origin extended community attribute, regardless if a site is connected to a single PE router or multiple PE routers. Configuring this attribute prevents routing loops from occurring when a site is multihomed. The SOO extended community attribute is configured on the interface and is propagated into BGP through redistribution. The SOO should not be configured for stub sites or sites that are not multihomed. **IP Extended Community-List Configuration Mode** Named and numbered extended community lists can be configured in IP Extended community-list configuration mode. To enter IP Extended community-list configuration mode, enter the **ip extcommunity-list** command with either the **expanded** or **standard** keyword followed by the extended community list name. This configuration mode supports all of the functions that are available in global configuration mode. In addition, you can perform the following operations: Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IRP-118	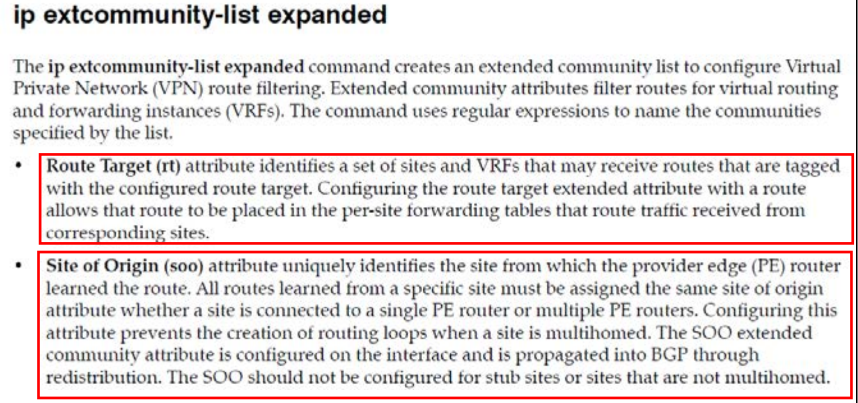 **ip extcommunity-list expanded** The **ip extcommunity-list expanded** command creates an extended community list to configure Virtual Private Network (VPN) route filtering. Extended community attributes filter routes for virtual routing and forwarding instances (VRFs). The command uses regular expressions to name the communities specified by the list. • **Route Target (rt)** attribute identifies a set of sites and VRFs that may receive routes that are tagged with the configured route target. Configuring the route target extended attribute with a route allows that route to be placed in the per-site forwarding tables that route traffic received from corresponding sites. • **Site of Origin (soo)** attribute uniquely identifies the site from which the provider edge (PE) router learned the route. All routes learned from a specific site must be assigned the same site of origin attribute whether a site is connected to a single PE router or multiple PE routers. Configuring this attribute prevents the creation of routing loops when a site is multihomed. The SOO extended community attribute is configured on the interface and is propagated into BGP through redistribution. The SOO should not be configured for stub sites or sites that are not multihomed. Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1540
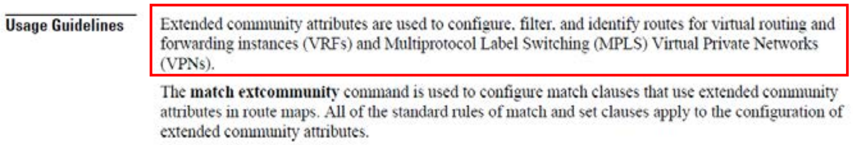 **Usage Guidelines** Extended community attributes are used to configure, filter, and identify routes for virtual routing and forwarding instances (VRFs) and Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). The **match extcommunity** command is used to configure match clauses that use extended community attributes in route maps. All of the standard rules of match and set clauses apply to the configuration of extended community attributes. Cisco IOS IP Routing: EIGRP Command Reference (2011), at 92	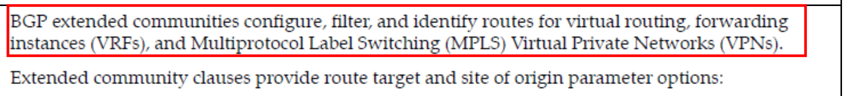 BGP extended communities configure, filter, and identify routes for virtual routing, forwarding instances (VRFs), and Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). Extended community clauses provide route target and site of origin parameter options: Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1502

Cisco	Arista
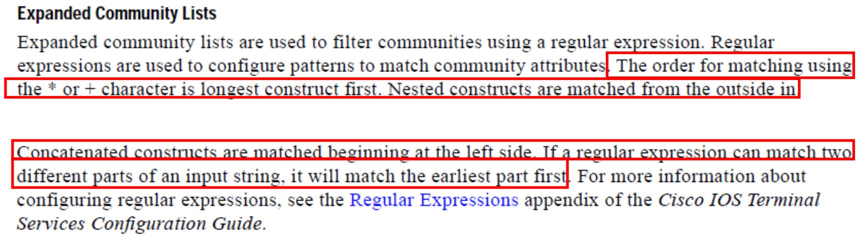 **Expanded Community Lists** Expanded community lists are used to filter communities using a regular expression. Regular expressions are used to configure patterns to match community attributes. The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it will match the earliest part first. For more information about configuring regular expressions, see the Regular Expressions appendix of the *Cisco IOS Terminal Services Configuration Guide*. Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IRP-113–14	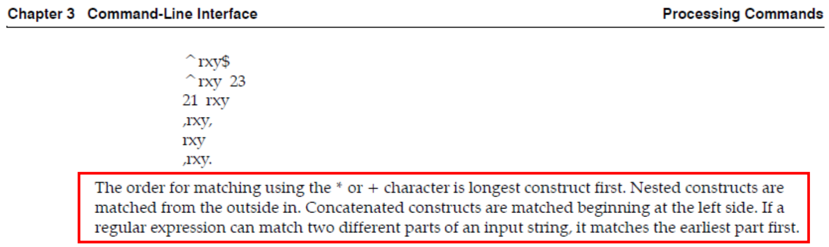 Chapter 3 Command-Line Interface Processing Commands ^rxy$ ^rxy 23 21 rxy ,rxy, rxy ,rxy. The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it matches the earliest part first. Arista Configuration Guide v. 4.13.6F (4/14/2014), at 105
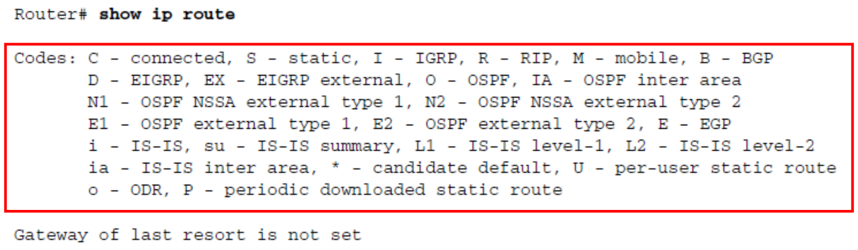 Router# **show ip route** Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 ia - IS-IS inter area, * - candidate default, U - per-user static route o - ODR, P - periodic downloaded static route Gateway of last resort is not set Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IP2R-553	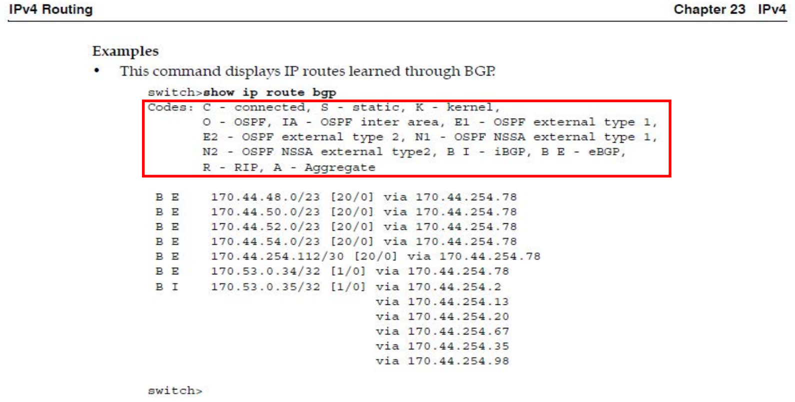 IPv4 Routing Chapter 23 IPv4 **Examples** • This command displays IP routes learned through BGP. switch>**show ip route bgp** Codes: C - connected, S - static, K - kernel, O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1, E2 - OSPF external type 2, N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type2, B I - iBGP, B E - eBGP, R - RIP, A - Aggregate B E 170.44.48.0/23 [20/0] via 170.44.254.78 B E 170.44.50.0/23 [20/0] via 170.44.254.78 B E 170.44.52.0/23 [20/0] via 170.44.254.78 B E 170.44.54.0/23 [20/0] via 170.44.254.78 B E 170.44.254.112/30 [20/0] via 170.44.254.78 B E 170.53.0.34/32 [1/0] via 170.44.254.78 B I 170.53.0.35/32 [1/0] via 170.44.254.2 via 170.44.254.13 via 170.44.254.20 via 170.44.254.67 via 170.44.254.35 via 170.44.254.98 switch> Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1188

Cisco	Arista
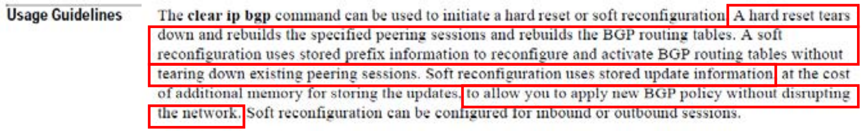 **Usage Guidelines** The **clear ip bgp** command can be used to initiate a hard reset or soft reconfiguration. A hard reset tears down and rebuilds the specified peering sessions and rebuilds the BGP routing tables. A soft reconfiguration uses stored prefix information to reconfigure and activate BGP routing tables without tearing down existing peering sessions. Soft reconfiguration uses stored update information at the cost of additional memory for storing the updates, to allow you to apply new BGP policy without disrupting the network. Soft reconfiguration can be configured for inbound or outbound sessions. Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IRP-69	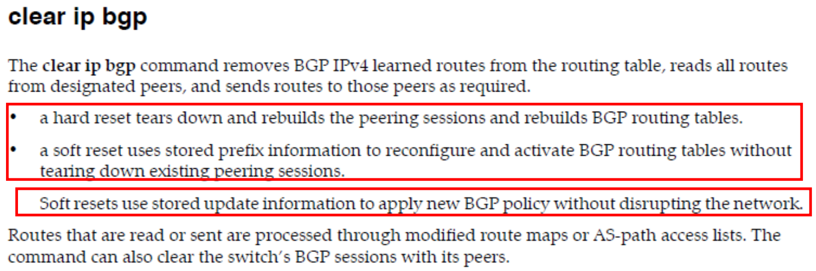 **clear ip bgp** The **clear ip bgp** command removes BGP IPv4 learned routes from the routing table, reads all routes from designated peers, and sends routes to those peers as required. • a hard reset tears down and rebuilds the peering sessions and rebuilds BGP routing tables. • a soft reset uses stored prefix information to reconfigure and activate BGP routing tables without tearing down existing peering sessions. Soft resets use stored update information to apply new BGP policy without disrupting the network. Routes that are read or sent are processed through modified route maps or AS-path access lists. The command can also clear the switch's BGP sessions with its peers. Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1527

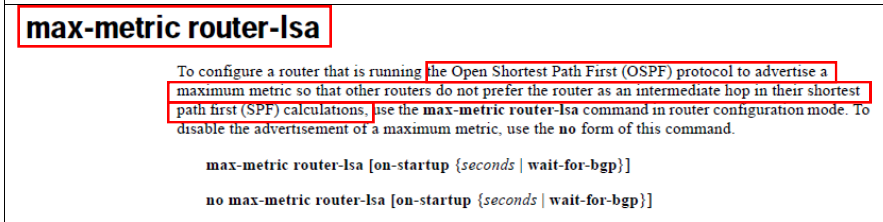
max-metric router-lsa
To configure a router that is running the Open Shortest Path First (OSPF) protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations, use the **max-metric router-lsa** command in router configuration mode. To disable the advertisement of a maximum metric, use the **no** form of this command.
max-metric router-lsa [on-startup {*seconds* | **wait-for-bgp**}**]**
no max-metric router-lsa [on-startup {*seconds* | **wait-for-bgp**}**]**
Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IP2R-591
	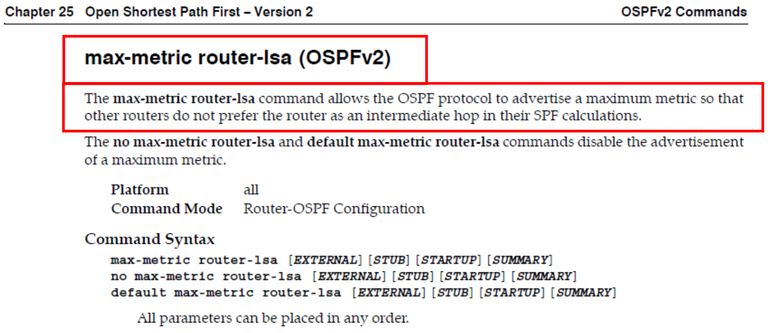
	Chapter 25 Open Shortest Path First – Version 2 OSPFv2 Commands
	max-metric router-lsa (OSPFv2)
	The **max-metric router-lsa** command allows the OSPF protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their SPF calculations.
	The **no max-metric router-lsa** and **default max-metric router-lsa** commands disable the advertisement of a maximum metric.
	Platform all
	Command Mode Router-OSPF Configuration
	Command Syntax
	`max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY]`
	`no max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY]`
	`default max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY]`
	All parameters can be placed in any order.
	Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1389

Cisco	Arista
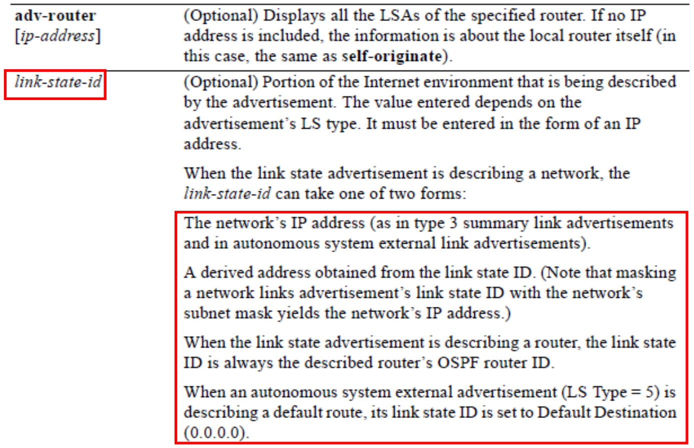 **adv-router** [*ip-address*] (Optional) Displays all the LSAs of the specified router. If no IP address is included, the information is about the local router itself (in this case, the same as **self-originate**). *link-state-id* (Optional) Portion of the Internet environment that is being described by the advertisement. The value entered depends on the advertisement's LS type. It must be entered in the form of an IP address. When the link state advertisement is describing a network, the *link-state-id* can take one of two forms: The network's IP address (as in type 3 summary link advertisements and in autonomous system external link advertisements). A derived address obtained from the link state ID. (Note that masking a network links advertisement's link state ID with the network's subnet mask yields the network's IP address.) When the link state advertisement is describing a router, the link state ID is always the described router's OSPF router ID. When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link state ID is set to Default Destination (0.0.0.0). Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IP2R-613	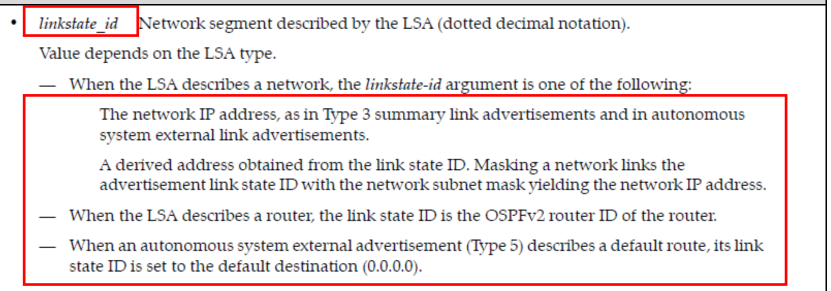 • *linkstate_id* Network segment described by the LSA (dotted decimal notation). Value depends on the LSA type. — When the LSA describes a network, the *linkstate-id* argument is one of the following: The network IP address, as in Type 3 summary link advertisements and in autonomous system external link advertisements. A derived address obtained from the link state ID. Masking a network links the advertisement link state ID with the network subnet mask yielding the network IP address. — When the LSA describes a router, the link state ID is the OSPFv2 router ID of the router. — When an autonomous system external advertisement (Type 5) describes a default route, its link state ID is set to the default destination (0.0.0.0). Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1404

Cisco

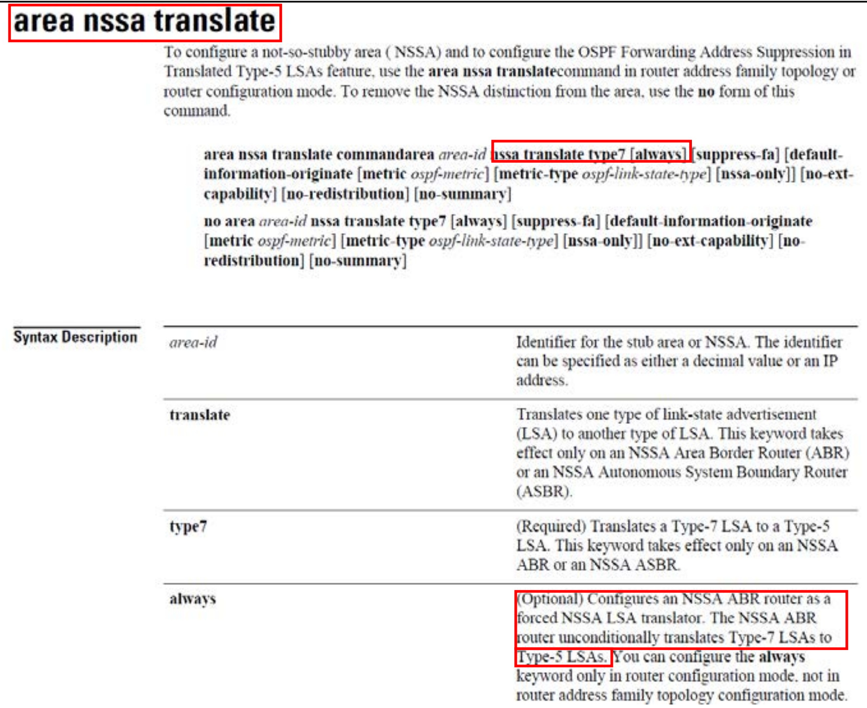

area nssa translate

To configure a not-so-stubby area (NSSA) and to configure the OSPF Forwarding Address Suppression in Translated Type-5 LSAs feature, use the **area nssa translate**command in router address family topology or router configuration mode. To remove the NSSA distinction from the area, use the **no** form of this command.

area nssa translate commandarea *area-id* **nssa translate type7 [always] [suppress-fa] [default-information-originate [metric** *ospf-metric*] **[metric-type** *ospf-link-state-type*] **[nssa-only]] [no-ext-capability] [no-redistribution] [no-summary]**

no area *area-id* **nssa translate type7 [always] [suppress-fa] [default-information-originate [metric** *ospf-metric*] **[metric-type** *ospf-link-state-type*] **[nssa-only]] [no-ext-capability] [no-redistribution] [no-summary]**

Syntax Description

area-id	Identifier for the stub area or NSSA. The identifier can be specified as either a decimal value or an IP address.
translate	Translates one type of link-state advertisement (LSA) to another type of LSA. This keyword takes effect only on an NSSA Area Border Router (ABR) or an NSSA Autonomous System Boundary Router (ASBR).
type7	(Required) Translates a Type-7 LSA to a Type-5 LSA. This keyword takes effect only on an NSSA ABR or an NSSA ASBR.
always	(Optional) Configures an NSSA ABR router as a forced NSSA LSA translator. The NSSA ABR router unconditionally translates Type-7 LSAs to Type-5 LSAs. You can configure the **always** keyword only in router configuration mode, not in router address family topology configuration mode.

Cisco IOS IP Routing: OSPF Command Reference (2011), at 15

Arista

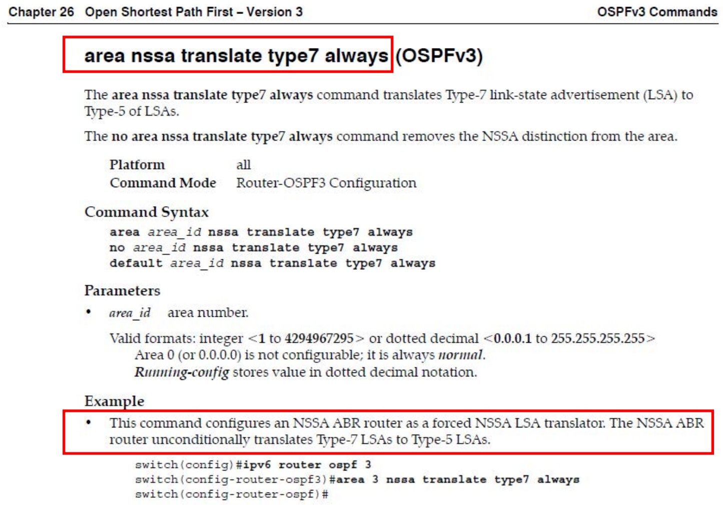

Chapter 26 Open Shortest Path First – Version 3 **OSPFv3 Commands**

area nssa translate type7 always (OSPFv3)

The **area nssa translate type7 always** command translates Type-7 link-state advertisement (LSA) to Type-5 of LSAs.

The **no area nssa translate type7 always** command removes the NSSA distinction from the area.

Platform	all
Command Mode	Router-OSPF3 Configuration

Command Syntax

```
area area_id nssa translate type7 always
no area_id nssa translate type7 always
default area_id nssa translate type7 always
```

Parameters

- *area_id* area number.

  Valid formats: integer <**1** to **4294967295**> or dotted decimal <**0.0.0.1** to **255.255.255.255**>
  Area 0 (or 0.0.0.0) is not configurable; it is always *normal*.
  Running-config stores value in dotted decimal notation.

Example

- This command configures an NSSA ABR router as a forced NSSA LSA translator. The NSSA ABR router unconditionally translates Type-7 LSAs to Type-5 LSAs.

```
switch(config)#ipv6 router ospf 3
switch(config-router-ospf3)#area 3 nssa translate type7 always
switch(config-router-ospf)#
```

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1451

Cisco	Arista
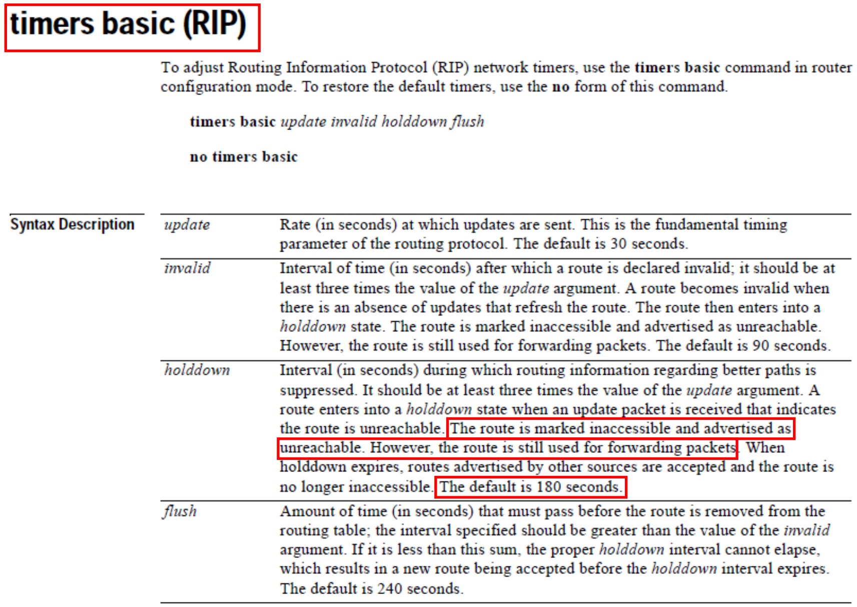 **timers basic (RIP)** To adjust Routing Information Protocol (RIP) network timers, use the **timers basic** command in router configuration mode. To restore the default timers, use the **no** form of this command. **timers basic** *update invalid holddown flush* **no timers basic** **Syntax Description** *update* — Rate (in seconds) at which updates are sent. This is the fundamental timing parameter of the routing protocol. The default is 30 seconds. *invalid* — Interval of time (in seconds) after which a route is declared invalid; it should be at least three times the value of the *update* argument. A route becomes invalid when there is an absence of updates that refresh the route. The route then enters into a *holddown* state. The route is marked inaccessible and advertised as unreachable. However, the route is still used for forwarding packets. The default is 90 seconds. *holddown* — Interval (in seconds) during which routing information regarding better paths is suppressed. It should be at least three times the value of the *update* argument. A route enters into a *holddown* state when an update packet is received that indicates the route is unreachable. The route is marked inaccessible and advertised as unreachable. However, the route is still used for forwarding packets. When holddown expires, routes advertised by other sources are accepted and the route is no longer inaccessible. The default is 180 seconds. *flush* — Amount of time (in seconds) that must pass before the route is removed from the routing table; the interval specified should be greater than the value of the *invalid* argument. If it is less than this sum, the proper *holddown* interval cannot elapse, which results in a new route being accepted before the *holddown* interval expires. The default is 240 seconds. Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IRP-811	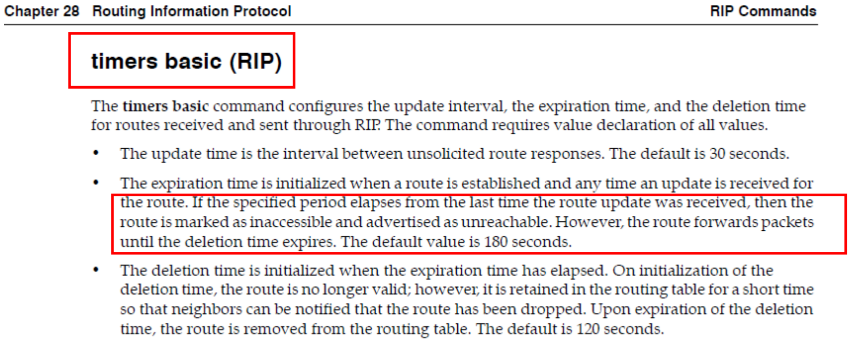 Chapter 28 Routing Information Protocol RIP Commands **timers basic (RIP)** The **timers basic** command configures the update interval, the expiration time, and the deletion time for routes received and sent through RIP. The command requires value declaration of all values. • The update time is the interval between unsolicited route responses. The default is 30 seconds. • The expiration time is initialized when a route is established and any time an update is received for the route. If the specified period elapses from the last time the route update was received, then the route is marked as inaccessible and advertised as unreachable. However, the route forwards packets until the deletion time expires. The default value is 180 seconds. • The deletion time is initialized when the expiration time has elapsed. On initialization of the deletion time, the route is no longer valid; however, it is retained in the routing table for a short time so that neighbors can be notified that the route has been dropped. Upon expiration of the deletion time, the route is removed from the routing table. The default is 120 seconds. Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1621

Cisco

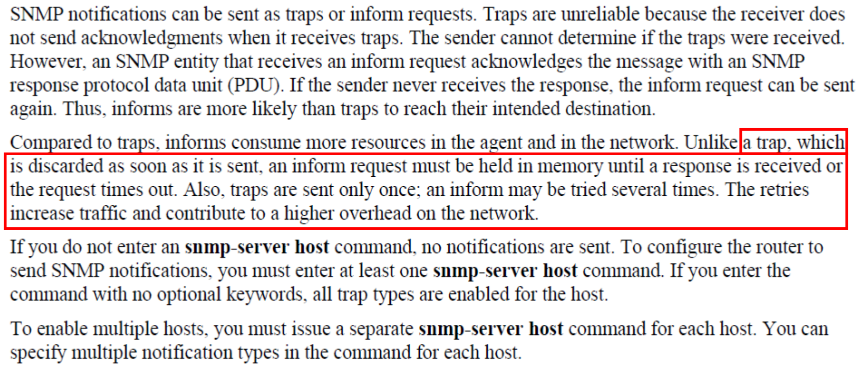

SNMP notifications can be sent as traps or inform requests. Traps are unreliable because the receiver does not send acknowledgments when it receives traps. The sender cannot determine if the traps were received. However, an SNMP entity that receives an inform request acknowledges the message with an SNMP response protocol data unit (PDU). If the sender never receives the response, the inform request can be sent again. Thus, informs are more likely than traps to reach their intended destination.

Compared to traps, informs consume more resources in the agent and in the network. Unlike a trap, which is discarded as soon as it is sent, an inform request must be held in memory until a response is received or the request times out. Also, traps are sent only once; an inform may be tried several times. The retries increase traffic and contribute to a higher overhead on the network.

If you do not enter an **snmp-server host** command, no notifications are sent. To configure the router to send SNMP notifications, you must enter at least one **snmp-server host** command. If you enter the command with no optional keywords, all trap types are enabled for the host.

To enable multiple hosts, you must issue a separate **snmp-server host** command for each host. You can specify multiple notification types in the command for each host.

Cisco IOS IP Switching Command Reference (2011), at 542

Arista

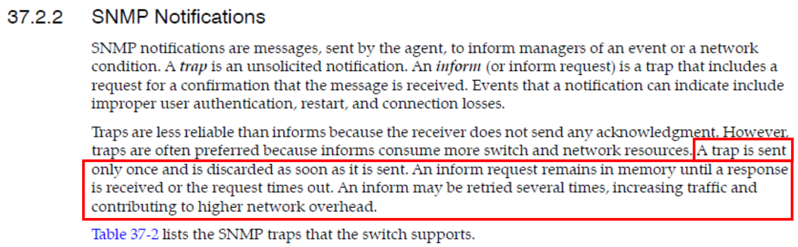

37.2.2 SNMP Notifications

SNMP notifications are messages, sent by the agent, to inform managers of an event or a network condition. A *trap* is an unsolicited notification. An *inform* (or inform request) is a trap that includes a request for a confirmation that the message is received. Events that a notification can indicate include improper user authentication, restart, and connection losses.

Traps are less reliable than informs because the receiver does not send any acknowledgment. However, traps are often preferred because informs consume more switch and network resources. A trap is sent only once and is discarded as soon as it is sent. An inform request remains in memory until a response is received or the request times out. An inform may be retried several times, increasing traffic and contributing to higher network overhead.

Table 37-2 lists the SNMP traps that the switch supports.

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1891

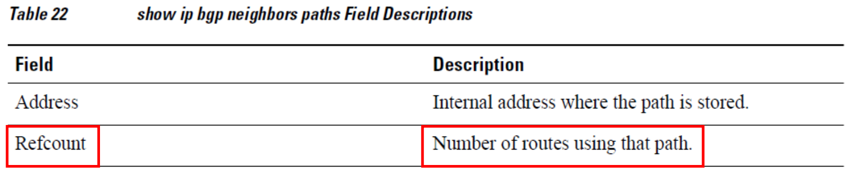

Table 22 show ip bgp neighbors paths Field Descriptions

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path.

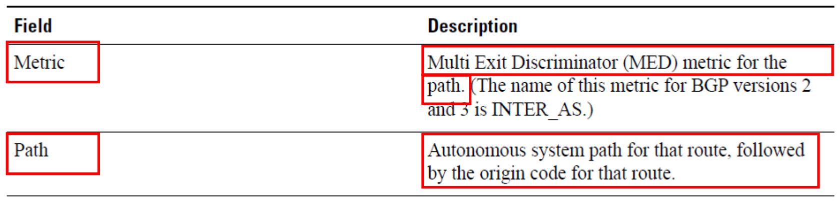

Field	Description
Metric	Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	Autonomous system path for that route, followed by the origin code for that route.

Cisco IOS Multiprotocol Label Switching Command Reference (2011), at 640-41

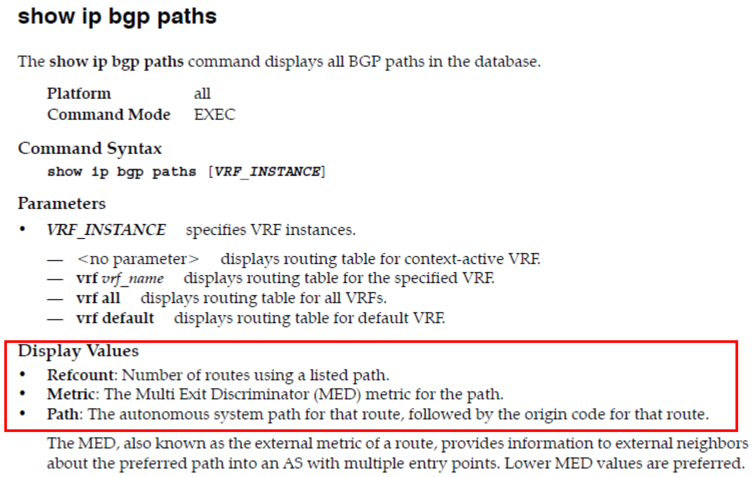

show ip bgp paths

The **show ip bgp paths** command displays all BGP paths in the database.

Platform all
Command Mode EXEC

Command Syntax

```
show ip bgp paths [VRF_INSTANCE]
```

Parameters

- *VRF_INSTANCE* specifies VRF instances.
  - — <no parameter> displays routing table for context-active VRF.
  - — **vrf** *vrf_name* displays routing table for the specified VRF.
  - — **vrf all** displays routing table for all VRFs.
  - — **vrf default** displays routing table for default VRF.

Display Values

- **Refcount**: Number of routes using a listed path.
- **Metric**: The Multi Exit Discriminator (MED) metric for the path.
- **Path**: The autonomous system path for that route, followed by the origin code for that route.

The MED, also known as the external metric of a route, provides information to external neighbors about the preferred path into an AS with multiple entry points. Lower MED values are preferred.

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1588

<table>
<tr><th>Cisco</th><th>Arista</th></tr>
<tr>
<td>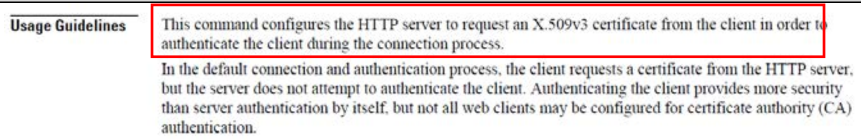
Usage Guidelines
This command configures the HTTP server to request an X.509v3 certificate from the client in order to authenticate the client during the connection process.
In the default connection and authentication process, the client requests a certificate from the HTTP server, but the server does not attempt to authenticate the client. Authenticating the client provides more security than server authentication by itself, but not all web clients may be configured for certificate authority (CA) authentication.

Cisco IOS HTTP Services Command Reference (2011), at 49</td>
<td>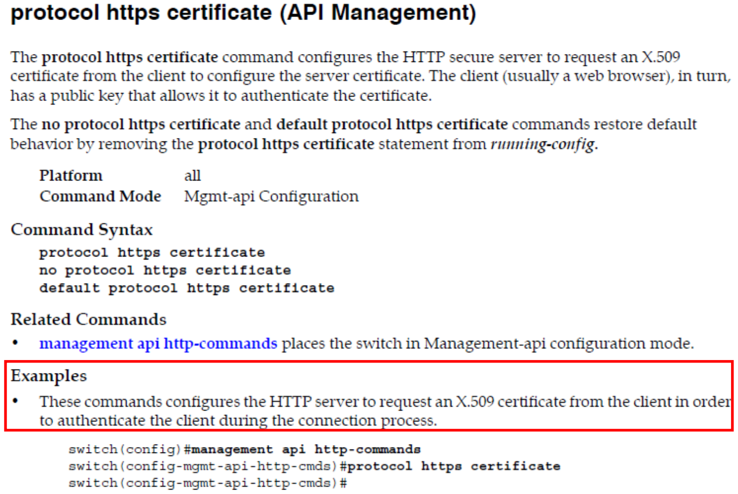
protocol https certificate (API Management)
The **protocol https certificate** command configures the HTTP secure server to request an X.509 certificate from the client to configure the server certificate. The client (usually a web browser), in turn, has a public key that allows it to authenticate the certificate.
The **no protocol https certificate** and **default protocol https certificate** commands restore default behavior by removing the **protocol https certificate** statement from ***running-config***.
Platform all
Command Mode Mgmt-api Configuration
Command Syntax
`protocol https certificate`
`no protocol https certificate`
`default protocol https certificate`
Related Commands
• **management api http-commands** places the switch in Management-api configuration mode.
Examples
• These commands configures the HTTP server to request an X.509 certificate from the client in order to authenticate the client during the connection process.
`switch(config)#management api http-commands`
`switch(config-mgmt-api-http-cmds)#protocol https certificate`
`switch(config-mgmt-api-http-cmds)#`

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 85</td>
</tr>
<tr>
<td>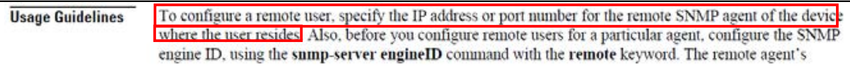
Usage Guidelines
To configure a remote user, specify the IP address or port number for the remote SNMP agent of the device where the user resides. Also, before you configure remote users for a particular agent, configure the SNMP engine ID, using the **snmp-server engineID** command with the **remote** keyword. The remote agent's

Cisco IOS SNMP Support Command Reference (2011), at 380</td>
<td>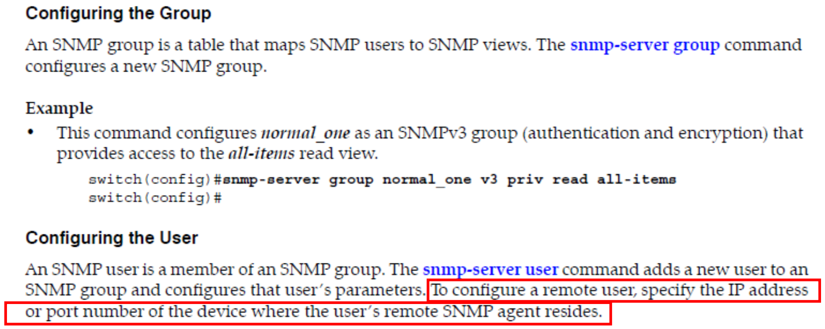
Configuring the Group
An SNMP group is a table that maps SNMP users to SNMP views. The **snmp-server group** command configures a new SNMP group.
Example
• This command configures *normal_one* as an SNMPv3 group (authentication and encryption) that provides access to the ***all-items*** read view.
`switch(config)#snmp-server group normal_one v3 priv read all-items`
`switch(config)#`
Configuring the User
An SNMP user is a member of an SNMP group. The **snmp-server user** command adds a new user to an SNMP group and configures that user's parameters. To configure a remote user, specify the IP address or port number of the device where the user's remote SNMP agent resides.

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1894</td>
</tr>
</table>

Cisco

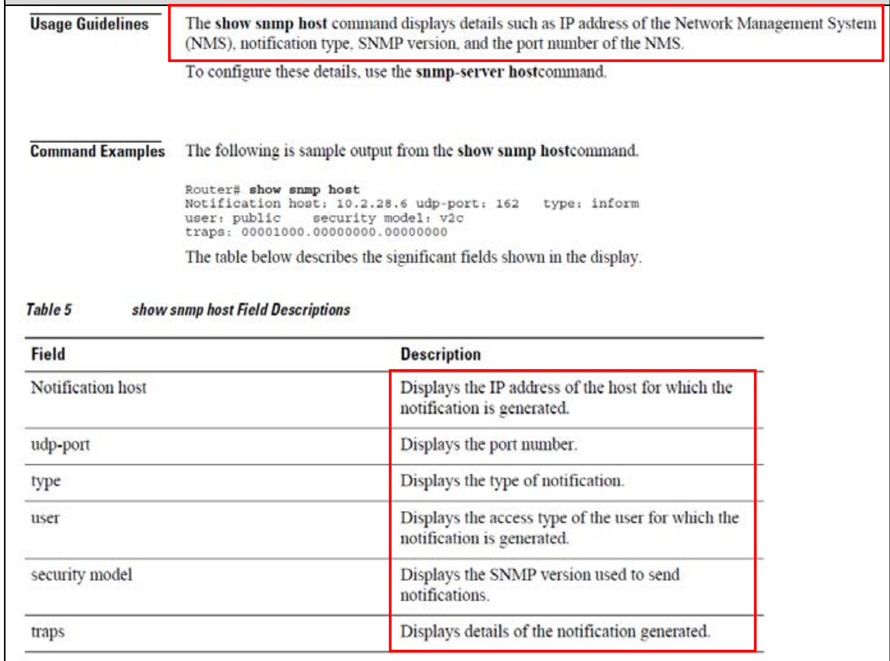

Usage Guidelines

The **show snmp host** command displays details such as IP address of the Network Management System (NMS), notification type, SNMP version, and the port number of the NMS.

To configure these details, use the **snmp-server host**command.

Command Examples

The following is sample output from the **show snmp host**command.

```
Router# show snmp host
Notification host: 10.2.28.6 udp-port: 162   type: inform
user: public    security model: v2c
traps: 00001000.00000000.00000000
```

The table below describes the significant fields shown in the display.

Table 5 *show snmp host Field Descriptions*

Field	Description
Notification host	Displays the IP address of the host for which the notification is generated.
udp-port	Displays the port number.
type	Displays the type of notification.
user	Displays the access type of the user for which the notification is generated.
security model	Displays the SNMP version used to send notifications.
traps	Displays details of the notification generated.

Cisco IOS SNMP Support Command Reference (July 2011), at 108–09

Arista

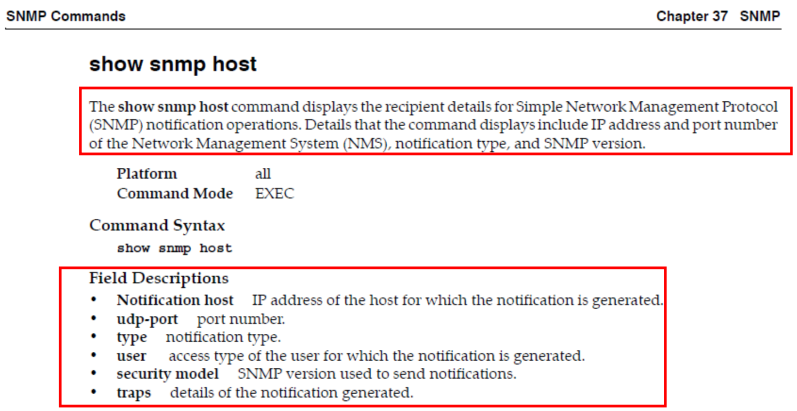

SNMP Commands Chapter 37 SNMP

show snmp host

The **show snmp host** command displays the recipient details for Simple Network Management Protocol (SNMP) notification operations. Details that the command displays include IP address and port number of the Network Management System (NMS), notification type, and SNMP version.

Platform all
Command Mode EXEC

Command Syntax

```
show snmp host
```

Field Descriptions

- **Notification host** IP address of the host for which the notification is generated.
- **udp-port** port number.
- **type** notification type.
- **user** access type of the user for which the notification is generated.
- **security model** SNMP version used to send notifications.
- **traps** details of the notification generated.

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1908

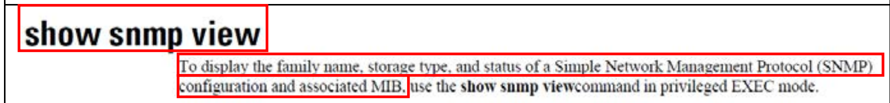

show snmp view

To display the family name, storage type, and status of a Simple Network Management Protocol (SNMP) configuration and associated MIB, use the **show snmp view**command in privileged EXEC mode.

Cisco IOS SNMP Support Command Reference (2011), at 140

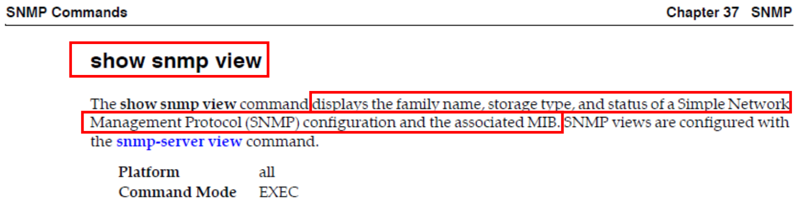

SNMP Commands Chapter 37 SNMP

show snmp view

The **show snmp view** command displays the family name, storage type, and status of a Simple Network Management Protocol (SNMP) configuration and the associated MIB. SNMP views are configured with the **snmp-server view** command.

Platform all
Command Mode EXEC

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1914

<table>
<tr><th>Cisco</th><th>Arista</th></tr>
<tr><td>

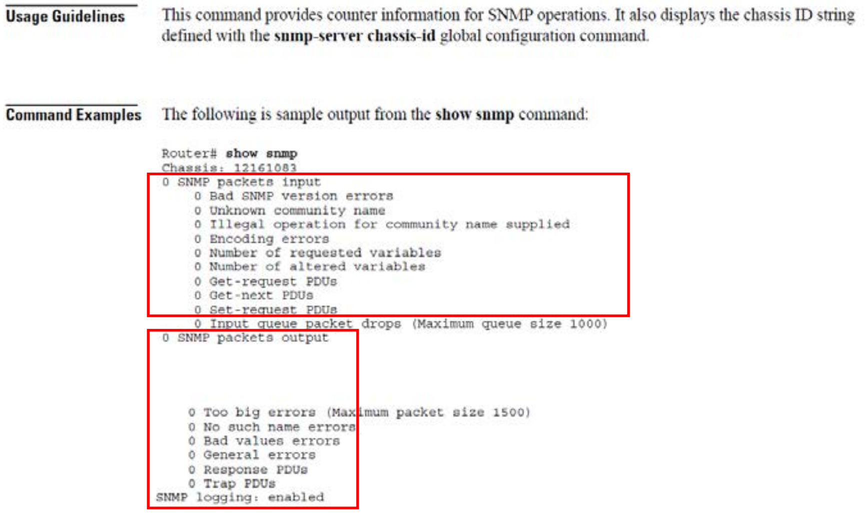

Usage Guidelines This command provides counter information for SNMP operations. It also displays the chassis ID string defined with the **snmp-server chassis-id** global configuration command.

Command Examples The following is sample output from the **show snmp** command:

```
Router# show snmp
Chassis: 12161083
0 SNMP packets input
    0 Bad SNMP version errors
    0 Unknown community name
    0 Illegal operation for community name supplied
    0 Encoding errors
    0 Number of requested variables
    0 Number of altered variables
    0 Get-request PDUs
    0 Get-next PDUs
    0 Set-request PDUs
    0 Input queue packet drops (Maximum queue size 1000)
0 SNMP packets output



   0 Too big errors (Maximum packet size 1500)
   0 No such name errors
   0 Bad values errors
   0 General errors
   0 Response PDUs
   0 Trap PDUs
SNMP logging: enabled
```

Cisco IOS SNMP Support Command Reference (2011), at 95-96

</td><td>

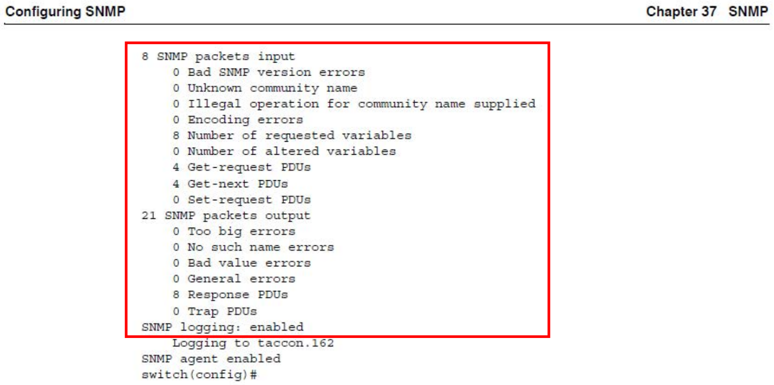

Configuring SNMP Chapter 37 SNMP

```
8 SNMP packets input
    0 Bad SNMP version errors
    0 Unknown community name
    0 Illegal operation for community name supplied
    0 Encoding errors
    8 Number of requested variables
    0 Number of altered variables
    4 Get-request PDUs
    4 Get-next PDUs
    0 Set-request PDUs
21 SNMP packets output
    0 Too big errors
    0 No such name errors
    0 Bad value errors
    0 General errors
    8 Response PDUs
    0 Trap PDUs
SNMP logging: enabled
    Logging to taccon.162
SNMP agent enabled
switch(config)#
```

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1896

</td></tr>
<tr><td>

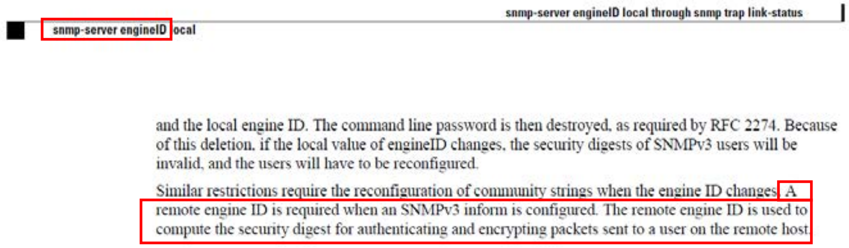

snmp-server engineID local through snmp trap link-status

snmp-server engineID local

and the local engine ID. The command line password is then destroyed, as required by RFC 2274. Because of this deletion, if the local value of engineID changes, the security digests of SNMPv3 users will be invalid, and the users will have to be reconfigured.

Similar restrictions require the reconfiguration of community strings when the engine ID changes. A remote engine ID is required when an SNMPv3 inform is configured. The remote engine ID is used to compute the security digest for authenticating and encrypting packets sent to a user on the remote host.

Cisco IOS SNMP Support Command Reference (2011), at 324

</td><td>

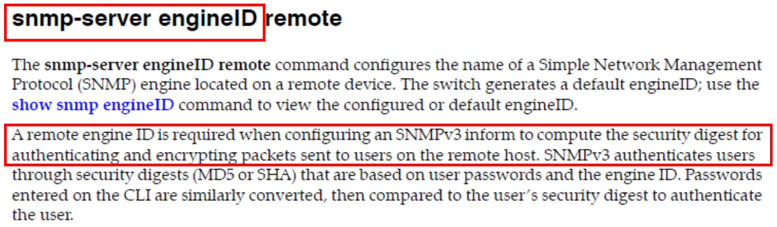

snmp-server engineID remote

The **snmp-server engineID remote** command configures the name of a Simple Network Management Protocol (SNMP) engine located on a remote device. The switch generates a default engineID; use the show snmp engineID command to view the configured or default engineID.

A remote engine ID is required when configuring an SNMPv3 inform to compute the security digest for authenticating and encrypting packets sent to users on the remote host. SNMPv3 authenticates users through security digests (MD5 or SHA) that are based on user passwords and the engine ID. Passwords entered on the CLI are similarly converted, then compared to the user's security digest to authenticate the user.

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 1920

</td></tr>
</table>

Cisco	Arista

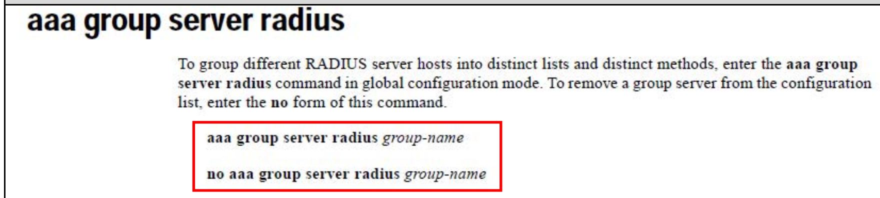
aaa group server radius
To group different RADIUS server hosts into distinct lists and distinct methods, enter the **aaa group server radius** command in global configuration mode. To remove a group server from the configuration list, enter the **no** form of this command.
aaa group server radius *group-name*
no aaa group server radius *group-name*
Cisco IOS Security Command Reference, Release 12.4 (2005), at SEC-74
	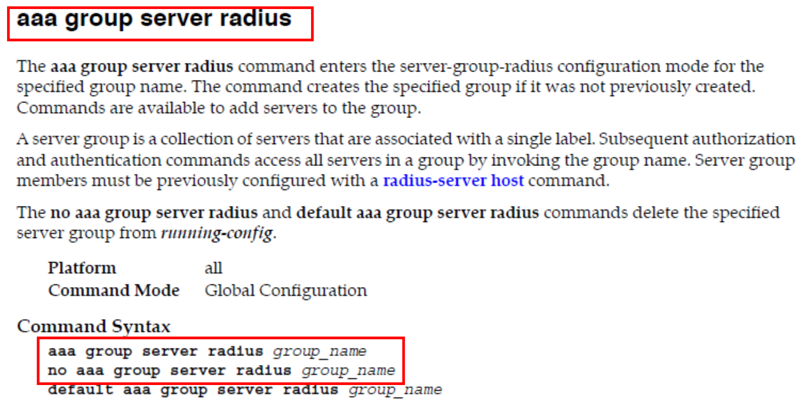
	aaa group server radius
	The **aaa group server radius** command enters the server-group-radius configuration mode for the specified group name. The command creates the specified group if it was not previously created. Commands are available to add servers to the group.
	A server group is a collection of servers that are associated with a single label. Subsequent authorization and authentication commands access all servers in a group by invoking the group name. Server group members must be previously configured with a **radius-server host** command.
	The **no aaa group server radius** and **default aaa group server radius** commands delete the specified server group from *running-config*.
	Platform all
	Command Mode Global Configuration
	Command Syntax
	`aaa group server radius group_name`
	`no aaa group server radius group_name`
	`default aaa group server radius group_name`
	Arista Configuration Guide v. 4.13.6F (4/14/2014), at 217

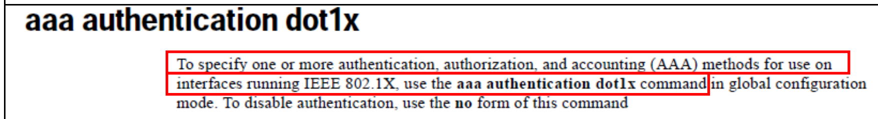
aaa authentication dot1x
To specify one or more authentication, authorization, and accounting (AAA) methods for use on interfaces running IEEE 802.1X, use the **aaa authentication dot1x** command in global configuration mode. To disable authentication, use the **no** form of this command
Cisco IOS Security Command Reference, Release 12.4 (2005), at SEC-32
	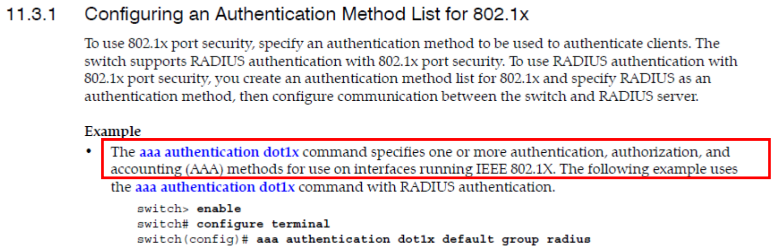
	11.3.1 Configuring an Authentication Method List for 802.1x
	To use 802.1x port security, specify an authentication method to be used to authenticate clients. The switch supports RADIUS authentication with 802.1x port security. To use RADIUS authentication with 802.1x port security, you create an authentication method list for 802.1x and specify RADIUS as an authentication method, then configure communication between the switch and RADIUS server.
	Example
	• The **aaa authentication dot1x** command specifies one or more authentication, authorization, and accounting (AAA) methods for use on interfaces running IEEE 802.1X. The following example uses the **aaa authentication dot1x** command with RADIUS authentication.
	`switch> enable`
	`switch# configure terminal`
	`switch(config)# aaa authentication dot1x default group radius`
	Arista Configuration Guide v. 4.13.6F (4/14/2014), at 551

Cisco

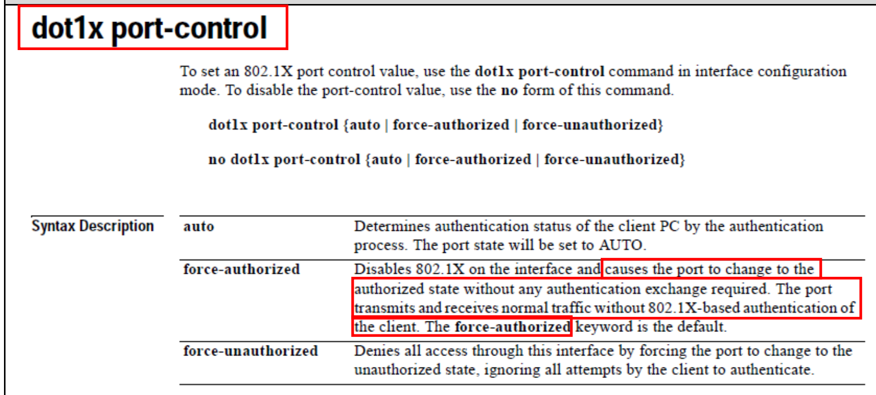

dot1x port-control

To set an 802.1X port control value, use the **dot1x port-control** command in interface configuration mode. To disable the port-control value, use the **no** form of this command.

dot1x port-control {**auto** | **force-authorized** | **force-unauthorized**}

no dot1x port-control {**auto** | **force-authorized** | **force-unauthorized**}

Syntax Description	**auto**	Determines authentication status of the client PC by the authentication process. The port state will be set to AUTO.
	force-authorized	Disables 802.1X on the interface and causes the port to change to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1X-based authentication of the client. The **force-authorized** keyword is the default.
	force-unauthorized	Denies all access through this interface by forcing the port to change to the unauthorized state, ignoring all attempts by the client to authenticate.

Cisco IOS Security Command Reference, Release 12.4 (2005), at SEC-457

Arista

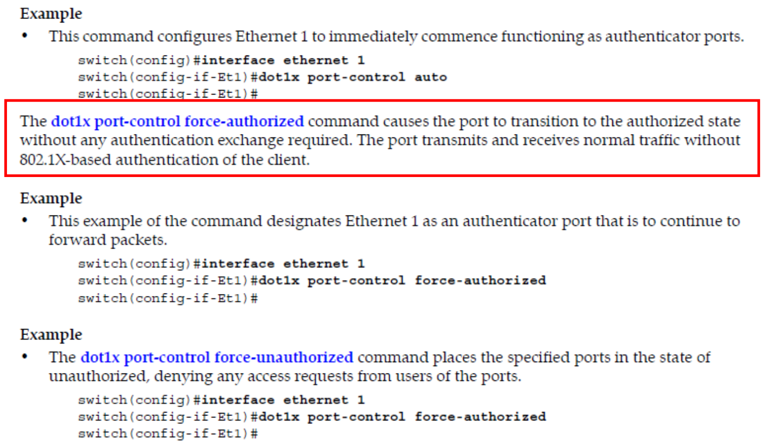

Example

- This command configures Ethernet 1 to immediately commence functioning as authenticator ports.

```
switch(config)#interface ethernet 1
switch(config-if-Et1)#dot1x port-control auto
switch(config-if-Et1)#
```

The **dot1x port-control force-authorized** command causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1X-based authentication of the client.

Example

- This example of the command designates Ethernet 1 as an authenticator port that is to continue to forward packets.

```
switch(config)#interface ethernet 1
switch(config-if-Et1)#dot1x port-control force-authorized
switch(config-if-Et1)#
```

Example

- The **dot1x port-control force-unauthorized** command places the specified ports in the state of unauthorized, denying any access requests from users of the ports.

```
switch(config)#interface ethernet 1
switch(config-if-Et1)#dot1x port-control force-authorized
switch(config-if-Et1)#
```

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 552

Cisco

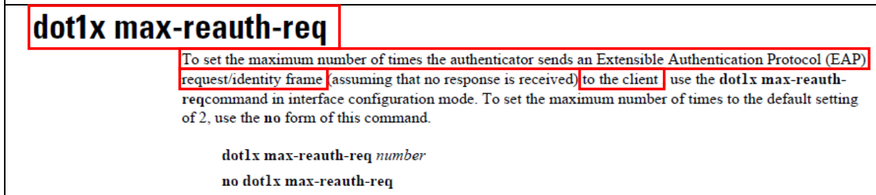

dot1x max-reauth-req

To set the maximum number of times the authenticator sends an Extensible Authentication Protocol (EAP) request/identity frame (assuming that no response is received) to the client, use the **dot1x max-reauth-reqcommand** in interface configuration mode. To set the maximum number of times to the default setting of 2, use the **no** form of this command.

dot1x max-reauth-req *number*

no dot1x max-reauth-req

Cisco IOS Security Command Reference: Commands D to L (2011), at 164

Arista

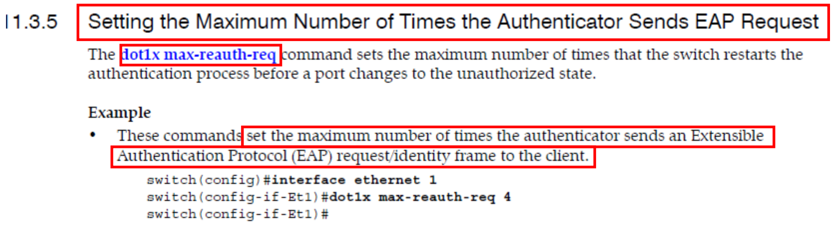

11.3.5 Setting the Maximum Number of Times the Authenticator Sends EAP Request

The **dot1x max-reauth-req** command sets the maximum number of times that the switch restarts the authentication process before a port changes to the unauthorized state.

Example

- These commands set the maximum number of times the authenticator sends an Extensible Authentication Protocol (EAP) request/identity frame to the client.

```
switch(config)#interface ethernet 1
switch(config-if-Et1)#dot1x max-reauth-req 4
switch(config-if-Et1)#
```

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 553

Cisco	Arista
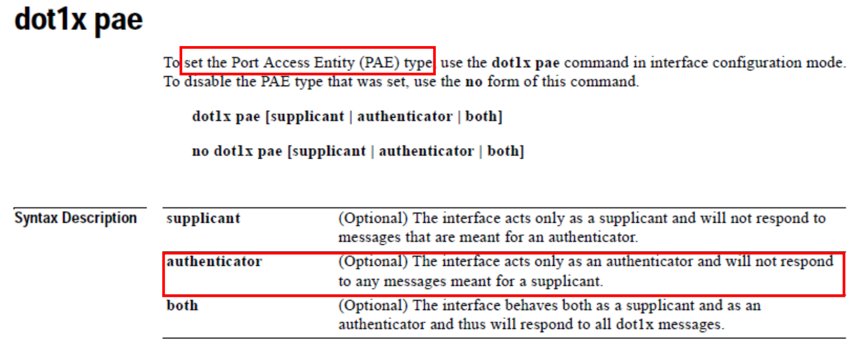 **dot1x pae** To set the Port Access Entity (PAE) type, use the **dot1x pae** command in interface configuration mode. To disable the PAE type that was set, use the **no** form of this command. **dot1x pae [supplicant | authenticator | both]** **no dot1x pae [supplicant | authenticator | both]** Syntax Description **supplicant** (Optional) The interface acts only as a supplicant and will not respond to messages that are meant for an authenticator. **authenticator** (Optional) The interface acts only as an authenticator and will not respond to any messages meant for a supplicant. **both** (Optional) The interface behaves both as a supplicant and as an authenticator and thus will respond to all dot1x messages. Cisco IOS Security Command Reference, Release 12.4 (2005), at SEC-456	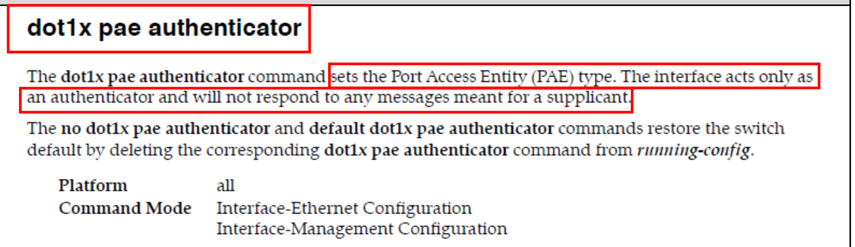 **dot1x pae authenticator** The **dot1x pae authenticator** command sets the Port Access Entity (PAE) type. The interface acts only as an authenticator and will not respond to any messages meant for a supplicant. The **no dot1x pae authenticator** and **default dot1x pae authenticator** commands restore the switch default by deleting the corresponding **dot1x pae authenticator** command from ***running-config***. **Platform** all **Command Mode** Interface-Ethernet Configuration Interface-Management Configuration Arista Configuration Guide v. 4.13.6F (4/14/2014), at 560
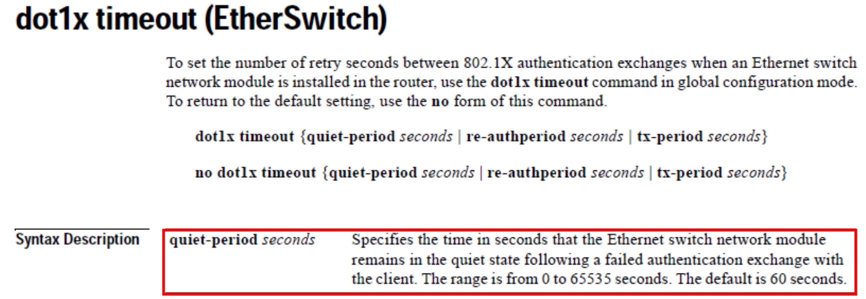 **dot1x timeout (EtherSwitch)** To set the number of retry seconds between 802.1X authentication exchanges when an Ethernet switch network module is installed in the router, use the **dot1x timeout** command in global configuration mode. To return to the default setting, use the **no** form of this command. **dot1x timeout** {**quiet-period** *seconds* | **re-authperiod** *seconds* | **tx-period** *seconds*} **no dot1x timeout** {**quiet-period** *seconds* | **re-authperiod** *seconds* | **tx-period** *seconds*} Syntax Description **quiet-period** *seconds* Specifies the time in seconds that the Ethernet switch network module remains in the quiet state following a failed authentication exchange with the client. The range is from 0 to 65535 seconds. The default is 60 seconds. Cisco IOS Security Command Reference, Release 12.4 (2005), at SEC-466	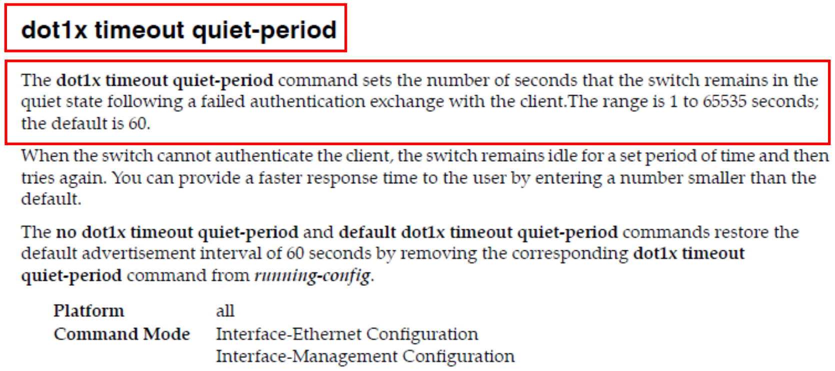 **dot1x timeout quiet-period** The **dot1x timeout quiet-period** command sets the number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client.The range is 1 to 65535 seconds; the default is 60. When the switch cannot authenticate the client, the switch remains idle for a set period of time and then tries again. You can provide a faster response time to the user by entering a number smaller than the default. The **no dot1x timeout quiet-period** and **default dot1x timeout quiet-period** commands restore the default advertisement interval of 60 seconds by removing the corresponding **dot1x timeout quiet-period** command from ***running-config***. **Platform** all **Command Mode** Interface-Ethernet Configuration Interface-Management Configuration Arista Configuration Guide v. 4.13.6F (4/14/2014), at 563

<table>
<tr><th>Cisco</th><th>Arista</th></tr>
<tr><td>

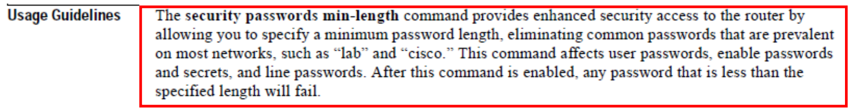

Usage Guidelines The **security passwords min-length** command provides enhanced security access to the router by allowing you to specify a minimum password length, eliminating common passwords that are prevalent on most networks, such as "lab" and "cisco." This command affects user passwords, enable passwords and secrets, and line passwords. After this command is enabled, any password that is less than the specified length will fail.

Cisco IOS Security Command Reference, Release 12.4 (2005), at SEC-943

</td><td>

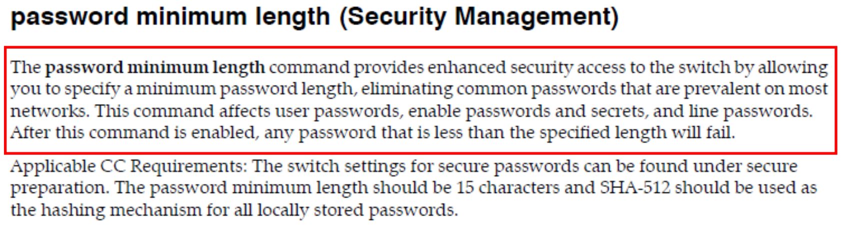

password minimum length (Security Management)

The **password minimum length** command provides enhanced security access to the switch by allowing you to specify a minimum password length, eliminating common passwords that are prevalent on most networks. This command affects user passwords, enable passwords and secrets, and line passwords. After this command is enabled, any password that is less than the specified length will fail.

Applicable CC Requirements: The switch settings for secure passwords can be found under secure preparation. The password minimum length should be 15 characters and SHA-512 should be used as the hashing mechanism for all locally stored passwords.

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 152

</td></tr>
<tr><td>

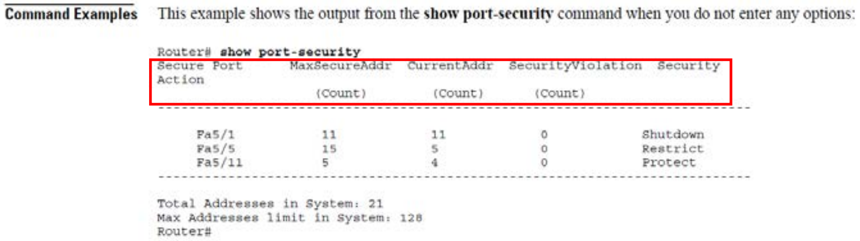

Command Examples This example shows the output from the **show port-security** command when you do not enter any options:

```
Router# show port-security
Secure Port      MaxSecureAddr  CurrentAddr  SecurityViolation  Security
Action
                     (Count)         (Count)       (Count)
------------------------------------------------------------------------

     Fa5/1           11            11            0             Shutdown
     Fa5/5           15            5             0             Restrict
     Fa5/11          5             4             0             Protect
------------------------------------------------------------------------

Total Addresses in System: 21
Max Addresses limit in System: 128
Router#
```

Cisco IOS Security Command Reference Commands S to Z (July 2011), at 692

</td><td>

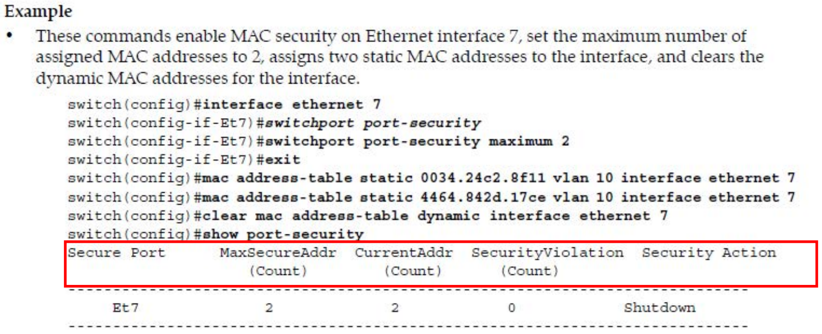

Example

- These commands enable MAC security on Ethernet interface 7, set the maximum number of assigned MAC addresses to 2, assigns two static MAC addresses to the interface, and clears the dynamic MAC addresses for the interface.

```
switch(config)#interface ethernet 7
switch(config-if-Et7)#switchport port-security
switch(config-if-Et7)#switchport port-security maximum 2
switch(config-if-Et7)#exit
switch(config)#mac address-table static 0034.24c2.8f11 vlan 10 interface ethernet 7
switch(config)#mac address-table static 4464.842d.17ce vlan 10 interface ethernet 7
switch(config)#clear mac address-table dynamic interface ethernet 7
switch(config)#show port-security
Secure Port      MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
                    (Count)        (Count)       (Count)
-----------------------------------------------------------------------------
     Et7              2              2              0              Shutdown
-----------------------------------------------------------------------------
```

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 624

</td></tr>
</table>

Cisco	Arista

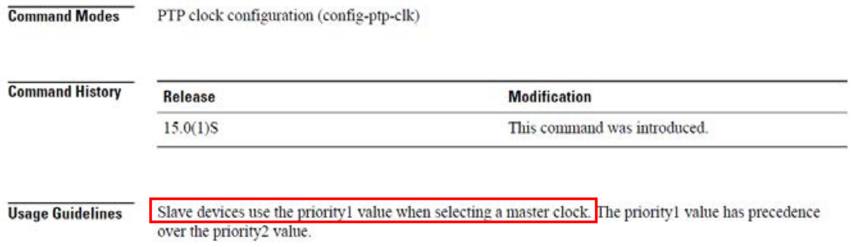
Command Modes PTP clock configuration (config-ptp-clk)
Command History
Release **Modification**
15.0(1)S This command was introduced.
Usage Guidelines Slave devices use the priority1 value when selecting a master clock. The priority1 value has precedence over the priority2 value.
Cisco IOS Interface and Hardware Component Command Reference (2011), at 1018
	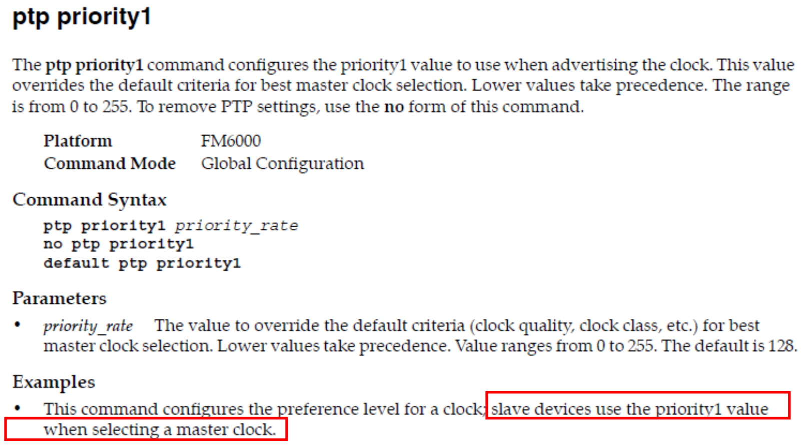
	ptp priority1
	The **ptp priority1** command configures the priority1 value to use when advertising the clock. This value overrides the default criteria for best master clock selection. Lower values take precedence. The range is from 0 to 255. To remove PTP settings, use the **no** form of this command.
	Platform FM6000
	Command Mode Global Configuration
	Command Syntax
	`ptp priority1` *`priority_rate`*
	`no ptp priority1`
	`default ptp priority1`
	Parameters
	• *priority_rate* The value to override the default criteria (clock quality, clock class, etc.) for best master clock selection. Lower values take precedence. Value ranges from 0 to 255. The default is 128.
	Examples
	• This command configures the preference level for a clock; slave devices use the priority1 value when selecting a master clock.
	Arista Configuration Guide v. 4.13.6F (4/14/2014), at 318

Cisco	Arista
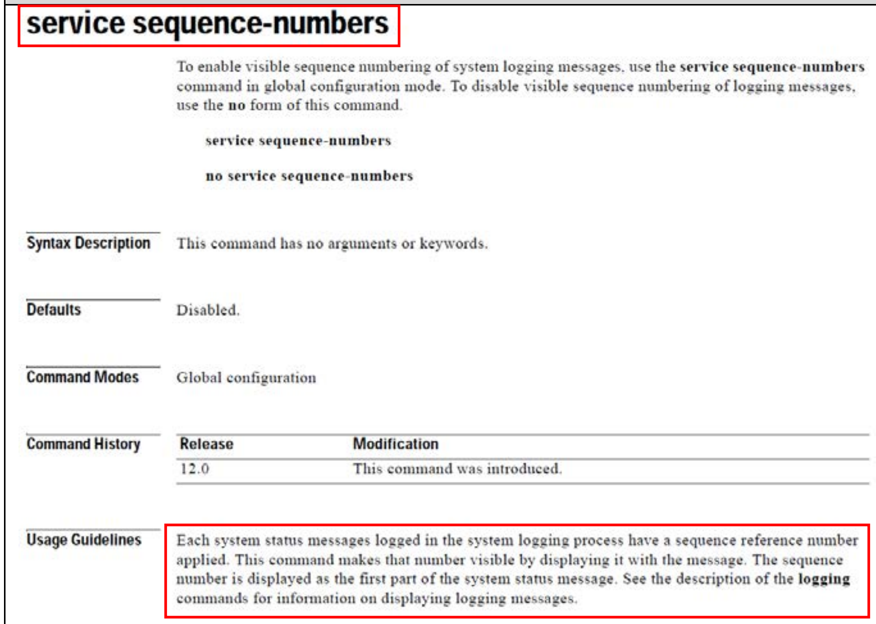 **service sequence-numbers** To enable visible sequence numbering of system logging messages, use the **service sequence-numbers** command in global configuration mode. To disable visible sequence numbering of logging messages, use the **no** form of this command. **service sequence-numbers** **no service sequence-numbers** **Syntax Description** This command has no arguments or keywords. **Defaults** Disabled. **Command Modes** Global configuration **Command History** **Release** **Modification** 12.0 This command was introduced. **Usage Guidelines** Each system status messages logged in the system logging process have a sequence reference number applied. This command makes that number visible by displaying it with the message. The sequence number is displayed as the first part of the system status message. See the description of the **logging** commands for information on displaying logging messages. Cisco IOS Configuration Fundamentals Command Reference Release 12.4T (2005), at CF-472	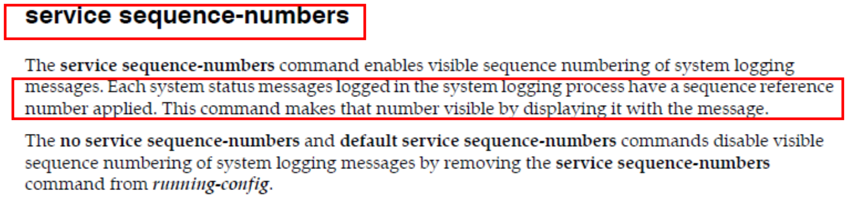 **service sequence-numbers** The **service sequence-numbers** command enables visible sequence numbering of system logging messages. Each system status messages logged in the system logging process have a sequence reference number applied. This command makes that number visible by displaying it with the message. The **no service sequence-numbers** and **default service sequence-numbers** commands disable visible sequence numbering of system logging messages by removing the **service sequence-numbers** command from *running-config*. Arista Configuration Guide v. 4.13.6F (4/14/2014), at 380

<table>
<tr><th>Cisco</th><th>Arista</th></tr>
<tr>
<td>
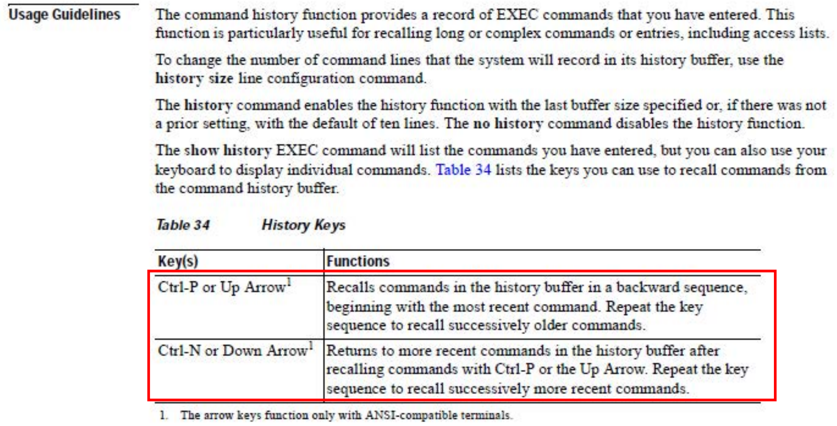

Usage Guidelines

The command history function provides a record of EXEC commands that you have entered. This function is particularly useful for recalling long or complex commands or entries, including access lists.

To change the number of command lines that the system will record in its history buffer, use the **history size** line configuration command.

The **history** command enables the history function with the last buffer size specified or, if there was not a prior setting, with the default of ten lines. The **no history** command disables the history function.

The **show history** EXEC command will list the commands you have entered, but you can also use your keyboard to display individual commands. Table 34 lists the keys you can use to recall commands from the command history buffer.

Table 34 History Keys

<table>
<tr><th>Key(s)</th><th>Functions</th></tr>
<tr><td>Ctrl-P or Up Arrow[1]</td><td>Recalls commands in the history buffer in a backward sequence, beginning with the most recent command. Repeat the key sequence to recall successively older commands.</td></tr>
<tr><td>Ctrl-N or Down Arrow[1]</td><td>Returns to more recent commands in the history buffer after recalling commands with Ctrl-P or the Up Arrow. Repeat the key sequence to recall successively more recent commands.</td></tr>
</table>

1. The arrow keys function only with ANSI-compatible terminals.

Cisco IOS Configuration Fundamentals Command Reference (2010), at CF-237
</td>
<td>
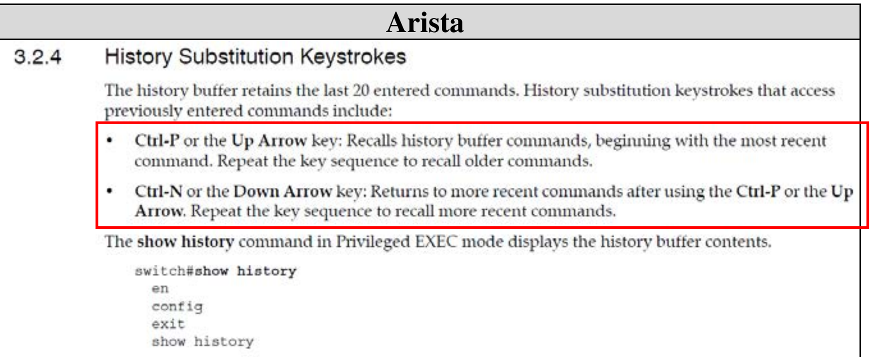

3.2.4 History Substitution Keystrokes

The history buffer retains the last 20 entered commands. History substitution keystrokes that access previously entered commands include:

- **Ctrl-P** or the **Up Arrow** key: Recalls history buffer commands, beginning with the most recent command. Repeat the key sequence to recall older commands.
- **Ctrl-N** or the **Down Arrow** key: Returns to more recent commands after using the **Ctrl-P** or the **Up Arrow**. Repeat the key sequence to recall more recent commands.

The **show history** command in Privileged EXEC mode displays the history buffer contents.

<pre>
switch#show history
  en
  config
  exit
  show history
</pre>

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 103
</td>
</tr>
<tr>
<td>
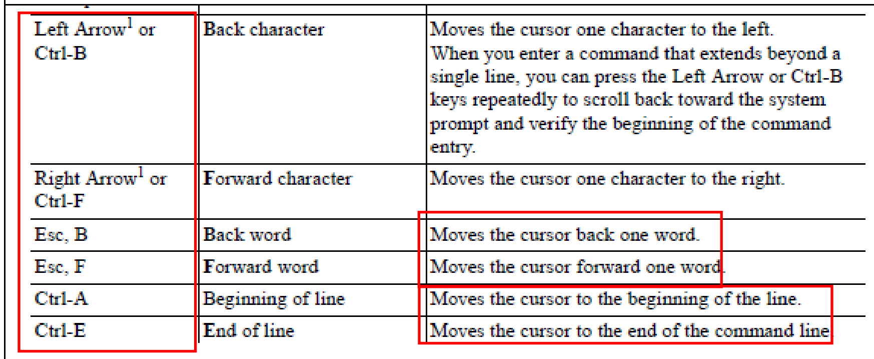

<table>
<tr><td>Left Arrow[1] or Ctrl-B</td><td>Back character</td><td>Moves the cursor one character to the left. When you enter a command that extends beyond a single line, you can press the Left Arrow or Ctrl-B keys repeatedly to scroll back toward the system prompt and verify the beginning of the command entry.</td></tr>
<tr><td>Right Arrow[1] or Ctrl-F</td><td>Forward character</td><td>Moves the cursor one character to the right.</td></tr>
<tr><td>Esc, B</td><td>Back word</td><td>Moves the cursor back one word.</td></tr>
<tr><td>Esc, F</td><td>Forward word</td><td>Moves the cursor forward one word.</td></tr>
<tr><td>Ctrl-A</td><td>Beginning of line</td><td>Moves the cursor to the beginning of the line.</td></tr>
<tr><td>Ctrl-E</td><td>End of line</td><td>Moves the cursor to the end of the command line</td></tr>
</table>

Cisco IOS Configuration Fundamentals Command Reference (2010), at CF-189
</td>
<td>
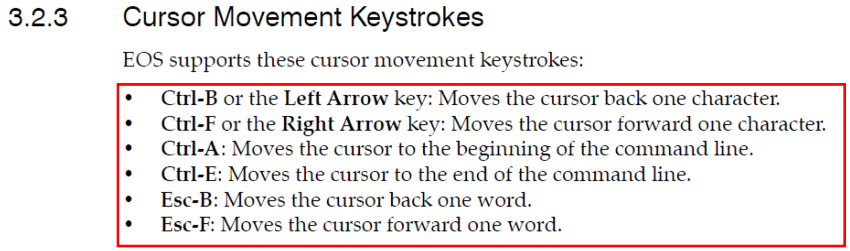

3.2.3 Cursor Movement Keystrokes

EOS supports these cursor movement keystrokes:

- **Ctrl-B** or the **Left Arrow** key: Moves the cursor back one character.
- **Ctrl-F** or the **Right Arrow** key: Moves the cursor forward one character.
- **Ctrl-A**: Moves the cursor to the beginning of the command line.
- **Ctrl-E**: Moves the cursor to the end of the command line.
- **Esc-B**: Moves the cursor back one word.
- **Esc-F**: Moves the cursor forward one word.

Arista Configuration Guide v. 4.13.6F (4/14/2014), at 102
</td>
</tr>
</table>

<table>
<tr><th>Cisco</th><th>Arista</th></tr>
<tr>
<td>
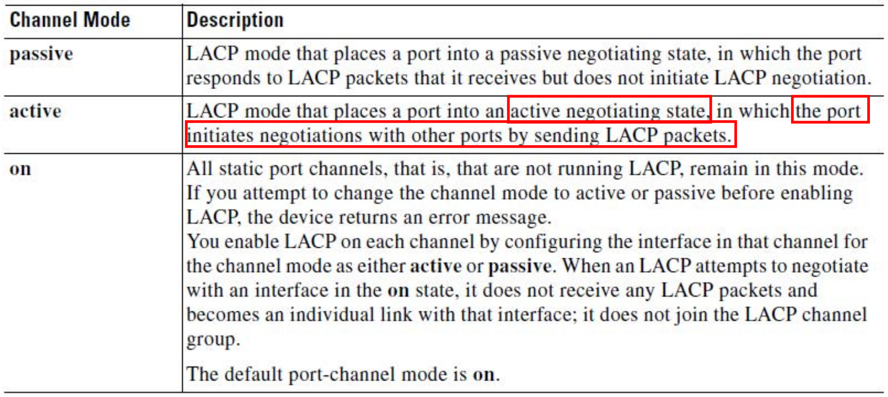
<table>
<tr><th>Channel Mode</th><th>Description</th></tr>
<tr><td><b>passive</b></td><td>LACP mode that places a port into a passive negotiating state, in which the port responds to LACP packets that it receives but does not initiate LACP negotiation.</td></tr>
<tr><td><b>active</b></td><td>LACP mode that places a port into an active negotiating state, in which the port initiates negotiations with other ports by sending LACP packets.</td></tr>
<tr><td><b>on</b></td><td>All static port channels, that is, that are not running LACP, remain in this mode. If you attempt to change the channel mode to active or passive before enabling LACP, the device returns an error message.
You enable LACP on each channel by configuring the interface in that channel for the channel mode as either <b>active</b> or <b>passive</b>. When an LACP attempts to negotiate with an interface in the <b>on</b> state, it does not receive any LACP packets and becomes an individual link with that interface; it does not join the LACP channel group.
The default port-channel mode is <b>on</b>.</td></tr>
</table>

Cisco NX-OS Interfaces Configuration Guide (2008), Release 4.0, at 5-9
</td>
<td>
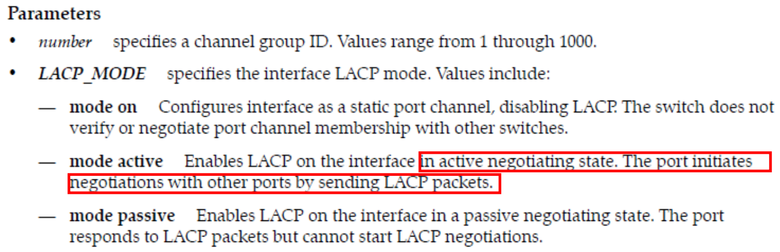
<b>Parameters</b>
• <i>number</i> specifies a channel group ID. Values range from 1 through 1000.
• <i>LACP_MODE</i> specifies the interface LACP mode. Values include:
— <b>mode on</b> Configures interface as a static port channel, disabling LACP. The switch does not verify or negotiate port channel membership with other switches.
— <b>mode active</b> Enables LACP on the interface in active negotiating state. The port initiates negotiations with other ports by sending LACP packets.
— <b>mode passive</b> Enables LACP on the interface in a passive negotiating state. The port responds to LACP packets but cannot start LACP negotiations.

Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 469
</td>
</tr>
<tr>
<td>
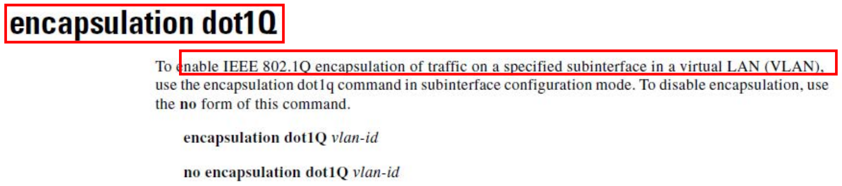
<b>encapsulation dot1Q</b>
To enable IEEE 802.1Q encapsulation of traffic on a specified subinterface in a virtual LAN (VLAN), use the encapsulation dot1q command in subinterface configuration mode. To disable encapsulation, use the <b>no</b> form of this command.
<b>encapsulation dot1Q</b> <i>vlan-id</i>
<b>no encapsulation dot1Q</b> <i>vlan-id</i>

Cisco NX-OS Interfaces Command Reference (2008), Release 4.0, at IF-8
</td>
<td>
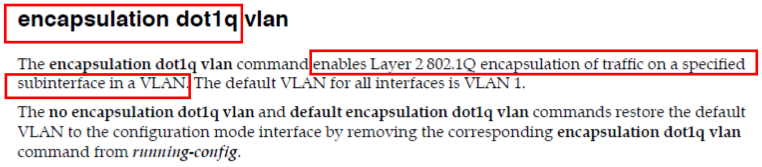
<b>encapsulation dot1q vlan</b>
The <b>encapsulation dot1q vlan</b> command enables Layer 2 802.1Q encapsulation of traffic on a specified subinterface in a VLAN. The default VLAN for all interfaces is VLAN 1.
The <b>no encapsulation dot1q vlan</b> and <b>default encapsulation dot1q vlan</b> commands restore the default VLAN to the configuration mode interface by removing the corresponding <b>encapsulation dot1q vlan</b> command from <i><b>running-config</b></i>.

Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 774
</td>
</tr>
</table>

Cisco	Arista

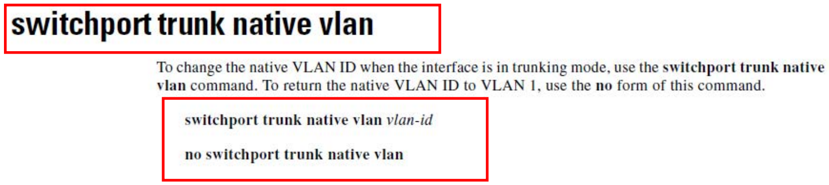
switchport trunk native vlan
To change the native VLAN ID when the interface is in trunking mode, use the **switchport trunk native vlan** command. To return the native VLAN ID to VLAN 1, use the **no** form of this command.
`switchport trunk native vlan` *`vlan-id`*
`no switchport trunk native vlan`
Cisco NX-OS Interfaces Command Reference (2008), Release 4.0, at IF-35
	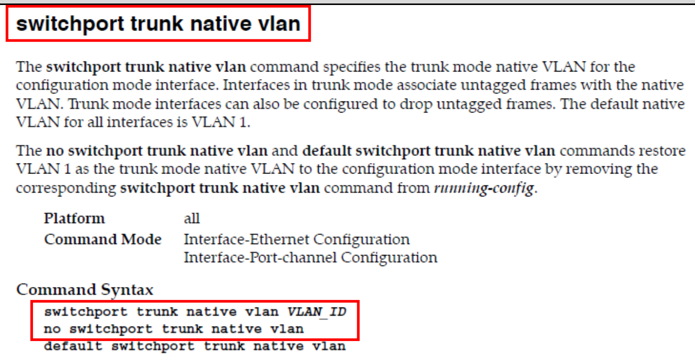
	switchport trunk native vlan
	The **switchport trunk native vlan** command specifies the trunk mode native VLAN for the configuration mode interface. Interfaces in trunk mode associate untagged frames with the native VLAN. Trunk mode interfaces can also be configured to drop untagged frames. The default native VLAN for all interfaces is VLAN 1.
	The **no switchport trunk native vlan** and **default switchport trunk native vlan** commands restore VLAN 1 as the trunk mode native VLAN to the configuration mode interface by removing the corresponding **switchport trunk native vlan** command from ***running-config***.
	Platform all
	Command Mode Interface-Ethernet Configuration
	Interface-Port-channel Configuration
	Command Syntax
	`switchport trunk native vlan` ***`VLAN_ID`***
	`no switchport trunk native vlan`
	`default switchport trunk native vlan`
	Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 800

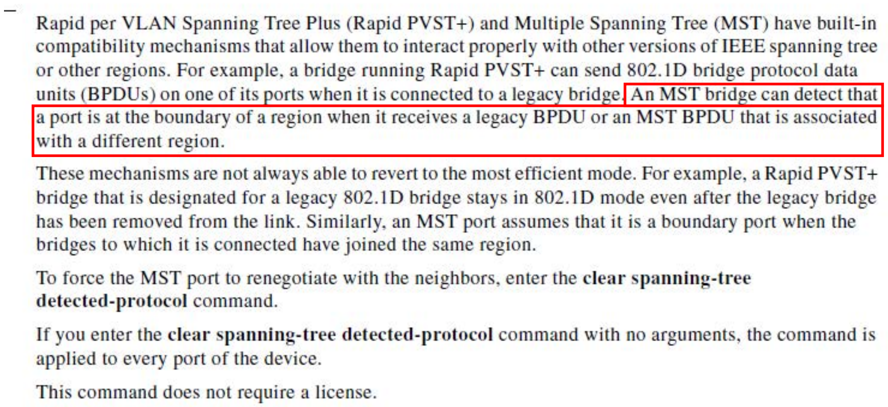 – Rapid per VLAN Spanning Tree Plus (Rapid PVST+) and Multiple Spanning Tree (MST) have built-in compatibility mechanisms that allow them to interact properly with other versions of IEEE spanning tree or other regions. For example, a bridge running Rapid PVST+ can send 802.1D bridge protocol data units (BPDUs) on one of its ports when it is connected to a legacy bridge. An MST bridge can detect that a port is at the boundary of a region when it receives a legacy BPDU or an MST BPDU that is associated with a different region. These mechanisms are not always able to revert to the most efficient mode. For example, a Rapid PVST+ bridge that is designated for a legacy 802.1D bridge stays in 802.1D mode even after the legacy bridge has been removed from the link. Similarly, an MST port assumes that it is a boundary port when the bridges to which it is connected have joined the same region. To force the MST port to renegotiate with the neighbors, enter the **clear spanning-tree detected-protocol** command. If you enter the **clear spanning-tree detected-protocol** command with no arguments, the command is applied to every port of the device. This command does not require a license. Cisco NX-OS Layer 2 Switching Command Reference (2008), Release 4.0, at L2-5	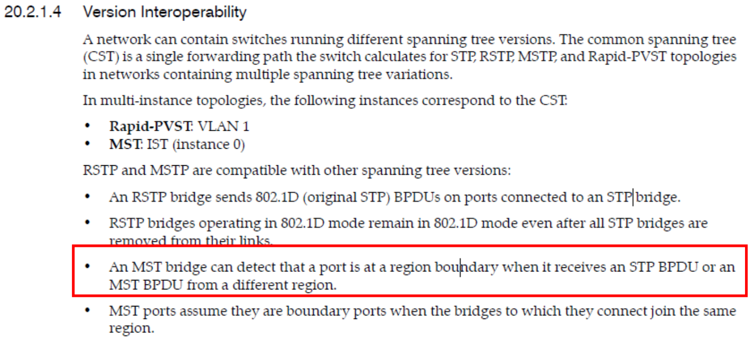 20.2.1.4 Version Interoperability A network can contain switches running different spanning tree versions. The common spanning tree (CST) is a single forwarding path the switch calculates for STP, RSTP, MSTP, and Rapid-PVST topologies in networks containing multiple spanning tree variations. In multi-instance topologies, the following instances correspond to the CST: • **Rapid-PVST**: VLAN 1 • **MST**: IST (instance 0) RSTP and MSTP are compatible with other spanning tree versions: • An RSTP bridge sends 802.1D (original STP) BPDUs on ports connected to an STP bridge. • RSTP bridges operating in 802.1D mode remain in 802.1D mode even after all STP bridges are removed from their links. • An MST bridge can detect that a port is at a region boundary when it receives an STP BPDU or an MST BPDU from a different region. • MST ports assume they are boundary ports when the bridges to which they connect join the same region. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 953

Cisco	Arista
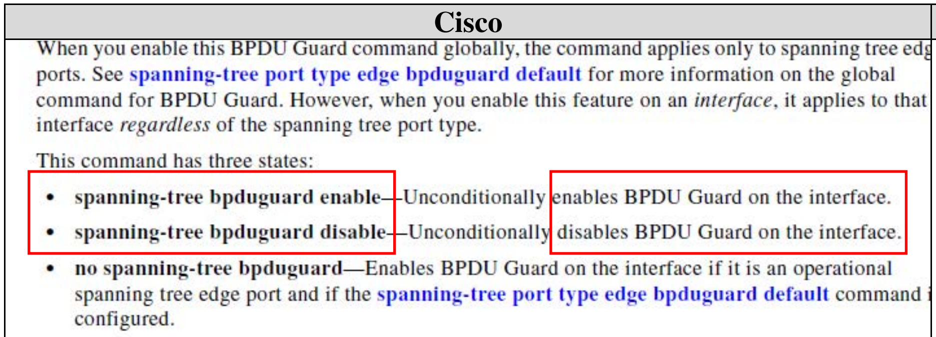 When you enable this BPDU Guard command globally, the command applies only to spanning tree edg ports. See **spanning-tree port type edge bpduguard default** for more information on the global command for BPDU Guard. However, when you enable this feature on an *interface*, it applies to that interface *regardless* of the spanning tree port type. This command has three states: • **spanning-tree bpduguard enable**—Unconditionally enables BPDU Guard on the interface. • **spanning-tree bpduguard disable**—Unconditionally disables BPDU Guard on the interface. • **no spanning-tree bpduguard**—Enables BPDU Guard on the interface if it is an operational spanning tree edge port and if the **spanning-tree port type edge bpduguard default** command i configured. Cisco NX-OS Layer 2 Switching Command Reference (2008), Release 4.0, at L2-31	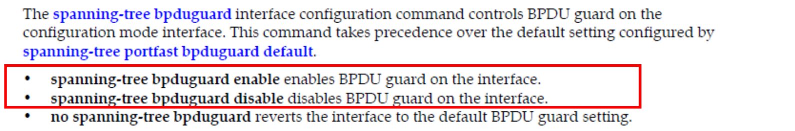 The **spanning-tree bpduguard** interface configuration command controls BPDU guard on the configuration mode interface. This command takes precedence over the default setting configured by **spanning-tree portfast bpduguard default**. • **spanning-tree bpduguard enable** enables BPDU guard on the interface. • **spanning-tree bpduguard disable** disables BPDU guard on the interface. • **no spanning-tree bpduguard** reverts the interface to the default BPDU guard setting. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 968
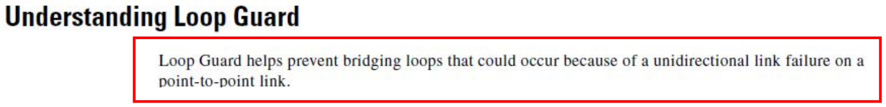 **Understanding Loop Guard** Loop Guard helps prevent bridging loops that could occur because of a unidirectional link failure on a point-to-point link. Cisco NX-OS Layer 2 Switching Configuration Guide (2008), Release 4.0, at 7-6	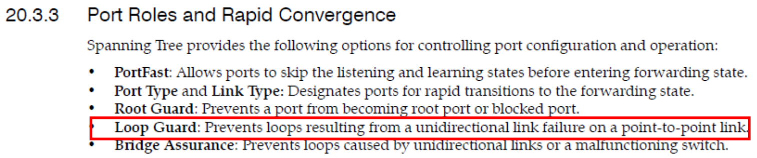 20.3.3 Port Roles and Rapid Convergence Spanning Tree provides the following options for controlling port configuration and operation: • **PortFast**: Allows ports to skip the listening and learning states before entering forwarding state. • **Port Type** and **Link Type:** Designates ports for rapid transitions to the forwarding state. • **Root Guard**: Prevents a port from becoming root port or blocked port. • **Loop Guard**: Prevents loops resulting from a unidirectional link failure on a point-to-point link. • **Bridge Assurance**: Prevents loops caused by unidirectional links or a malfunctioning switch. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 964
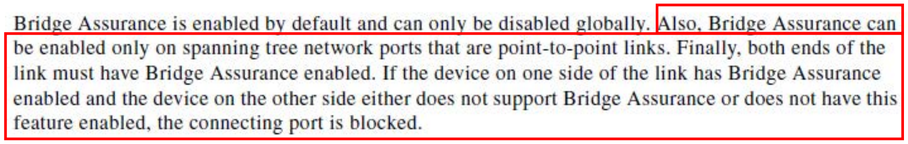 Bridge Assurance is enabled by default and can only be disabled globally. Also, Bridge Assurance can be enabled only on spanning tree network ports that are point-to-point links. Finally, both ends of the link must have Bridge Assurance enabled. If the device on one side of the link has Bridge Assurance enabled and the device on the other side either does not support Bridge Assurance or does not have this feature enabled, the connecting port is blocked. Cisco NX-OS Layer 2 Switching Configuration Guide (2008), Release 4.0, at 7-3	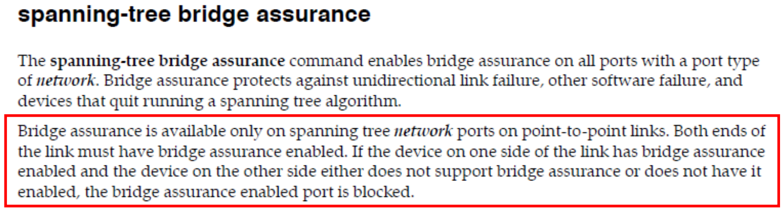 **spanning-tree bridge assurance** The **spanning-tree bridge assurance** command enables bridge assurance on all ports with a port type of ***network***. Bridge assurance protects against unidirectional link failure, other software failure, and devices that quit running a spanning tree algorithm. Bridge assurance is available only on spanning tree ***network*** ports on point-to-point links. Both ends of the link must have bridge assurance enabled. If the device on one side of the link has bridge assurance enabled and the device on the other side either does not support bridge assurance or does not have it enabled, the bridge assurance enabled port is blocked. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 1002
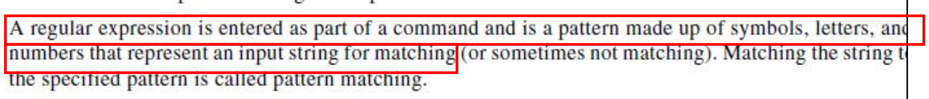 A regular expression is entered as part of a command and is a pattern made up of symbols, letters, anc numbers that represent an input string for matching (or sometimes not matching). Matching the string t the specified pattern is called pattern matching. Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at	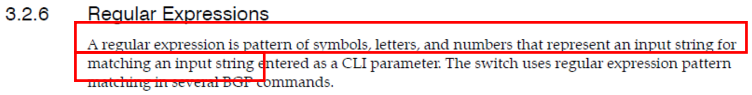 3.2.6 Regular Expressions A regular expression is pattern of symbols, letters, and numbers that represent an input string for matching an input string entered as a CLI parameter. The switch uses regular expression pattern matching in several BGP commands. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 106

<table>
<tr><th>Cisco</th><th>Arista</th></tr>
<tr><td>A-1</td><td></td></tr>
<tr><td>
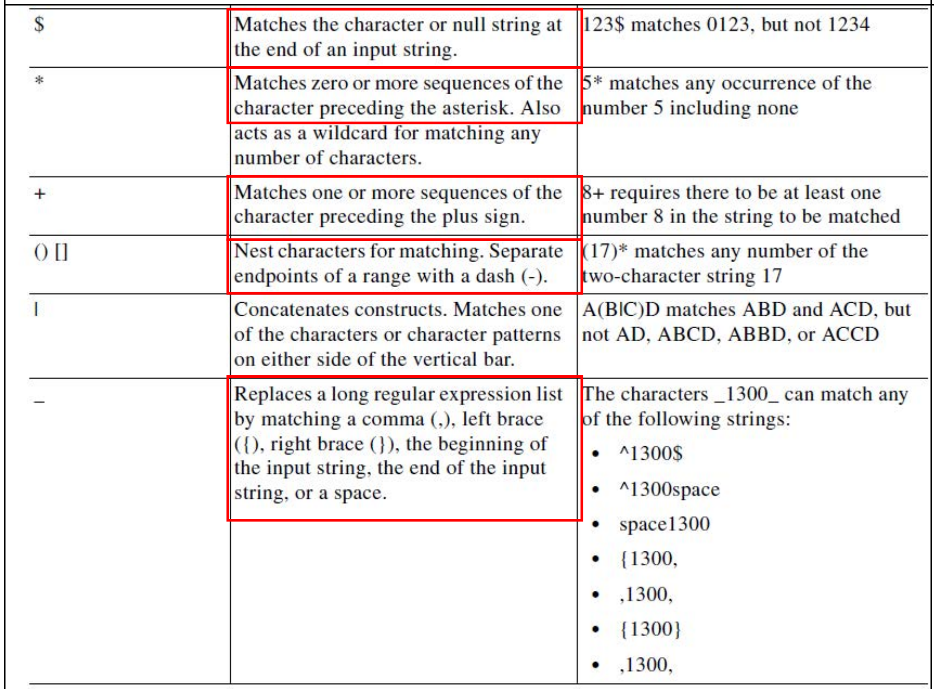
<table>
<tr><td>$</td><td>Matches the character or null string at the end of an input string.</td><td>123$ matches 0123, but not 1234</td></tr>
<tr><td>*</td><td>Matches zero or more sequences of the character preceding the asterisk. Also acts as a wildcard for matching any number of characters.</td><td>5* matches any occurrence of the number 5 including none</td></tr>
<tr><td>+</td><td>Matches one or more sequences of the character preceding the plus sign.</td><td>8+ requires there to be at least one number 8 in the string to be matched</td></tr>
<tr><td>() []</td><td>Nest characters for matching. Separate endpoints of a range with a dash (-).</td><td>(17)* matches any number of the two-character string 17</td></tr>
<tr><td>|</td><td>Concatenates constructs. Matches one of the characters or character patterns on either side of the vertical bar.</td><td>A(B|C)D matches ABD and ACD, but not AD, ABCD, ABBD, or ACCD</td></tr>
<tr><td>_</td><td>Replaces a long regular expression list by matching a comma (,), left brace ({), right brace (}), the beginning of the input string, the end of the input string, or a space.</td><td>The characters _1300_ can match any of the following strings:
• ^1300$
• ^1300space
• space1300
• {1300,
• ,1300,
• {1300}
• ,1300,</td></tr>
</table>

Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at A-2
</td><td>
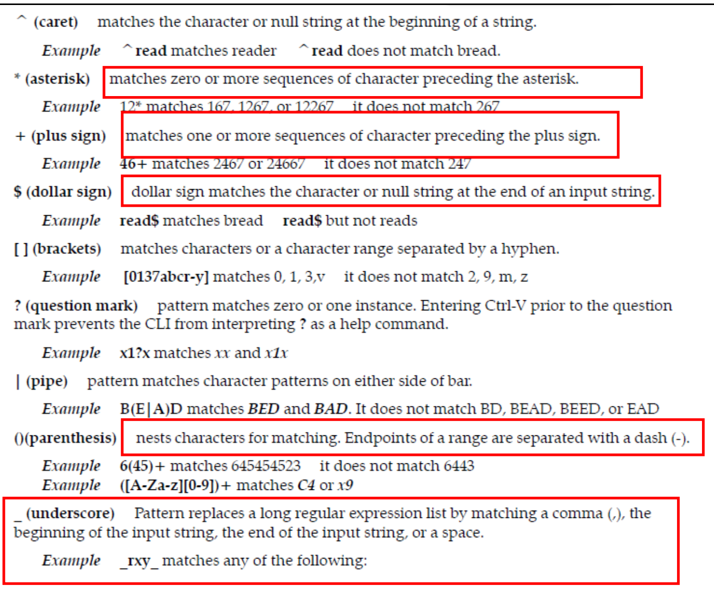
^ (caret) matches the character or null string at the beginning of a string.
Example ^**read** matches reader ^**read** does not match bread.
* **(asterisk)** matches zero or more sequences of character preceding the asterisk.
Example 12* matches 167, 1267, or 12267 it does not match 267
+ (plus sign) matches one or more sequences of character preceding the plus sign.
Example **46+** matches 2467 or 24667 it does not match 247
$ (dollar sign) dollar sign matches the character or null string at the end of an input string.
Example **read$** matches bread **read$** but not reads
[ ] (brackets) matches characters or a character range separated by a hyphen.
Example **[0137abcr-y]** matches 0, 1, 3,v it does not match 2, 9, m, z
? (question mark) pattern matches zero or one instance. Entering Ctrl-V prior to the question mark prevents the CLI from interpreting **?** as a help command.
Example **x1?x** matches *xx* and *x1x*
| (pipe) pattern matches character patterns on either side of bar.
Example B(E|A)D matches ***BED*** and ***BAD***. It does not match BD, BEAD, BEED, or EAD
()(parenthesis) nests characters for matching. Endpoints of a range are separated with a dash (-).
Example 6(45)+ matches 645454523 it does not match 6443
Example **([A-Za-z][0-9])+** matches *C4* or *x9*
_ (underscore) Pattern replaces a long regular expression list by matching a comma (,), the beginning of the input string, the end of the input string, or a space.
Example **_rxy_** matches any of the following:

Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 106
</td></tr>
<tr><td>
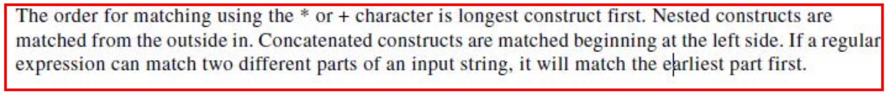
The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it will match the earliest part first.

Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at A-3
</td><td>
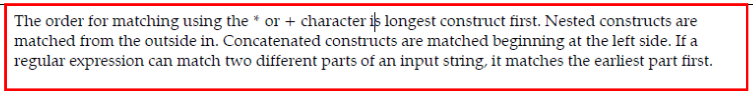
The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it matches the earliest part first.

Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 107
</td></tr>
</table>

Cisco	Arista
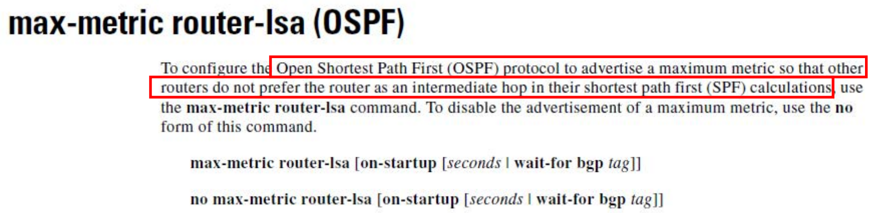 **max-metric router-lsa (OSPF)** To configure the Open Shortest Path First (OSPF) protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations, use the **max-metric router-lsa** command. To disable the advertisement of a maximum metric, use the **no** form of this command. **max-metric router-lsa** [**on-startup** [*seconds* | **wait-for bgp** *tag*]] **no max-metric router-lsa** [**on-startup** [*seconds* | **wait-for bgp** *tag*]] Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at L3-272	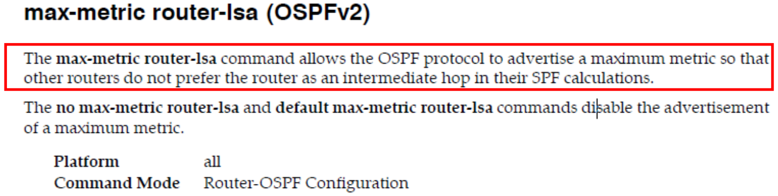 **max-metric router-lsa (OSPFv2)** The **max-metric router-lsa** command allows the OSPF protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their SPF calculations. The **no max-metric router-lsa** and **default max-metric router-lsa** commands disable the advertisement of a maximum metric. **Platform** all **Command Mode** Router-OSPF Configuration Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 1439
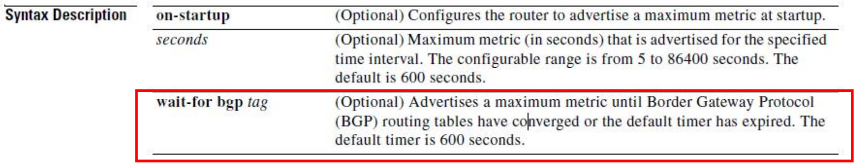 **Syntax Description** **on-startup** (Optional) Configures the router to advertise a maximum metric at startup. *seconds* (Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds. **wait-for bgp** *tag* (Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds. Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at L3-272	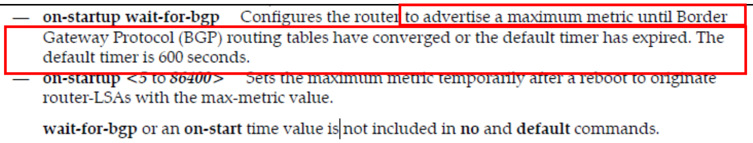 — **on-startup wait-for-bgp** Configures the router to advertise a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds. — **on-startup** <5 to *86400*> Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. **wait-for-bgp** or an **on-start** time value is not included in **no** and **default** commands. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 1439
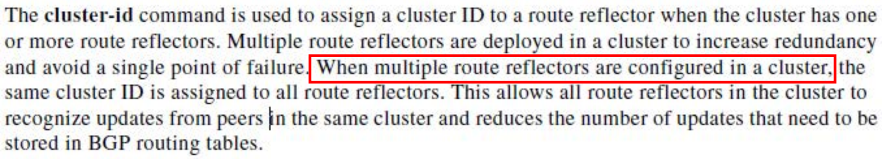 The **cluster-id** command is used to assign a cluster ID to a route reflector when the cluster has one or more route reflectors. Multiple route reflectors are deployed in a cluster to increase redundancy and avoid a single point of failure. When multiple route reflectors are configured in a cluster, the same cluster ID is assigned to all route reflectors. This allows all route reflectors in the cluster to recognize updates from peers in the same cluster and reduces the number of updates that need to be stored in BGP routing tables. Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at L3-564	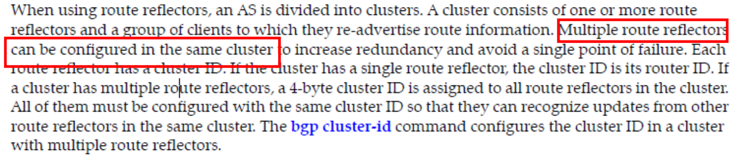 When using route reflectors, an AS is divided into clusters. A cluster consists of one or more route reflectors and a group of clients to which they re-advertise route information. Multiple route reflectors can be configured in the same cluster to increase redundancy and avoid a single point of failure. Each route reflector has a cluster ID. If the cluster has a single route reflector, the cluster ID is its router ID. If a cluster has multiple route reflectors, a 4-byte cluster ID is assigned to all route reflectors in the cluster. All of them must be configured with the same cluster ID so that they can recognize updates from other route reflectors in the same cluster. The **bgp cluster-id** command configures the cluster ID in a cluster with multiple route reflectors. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 1549

Cisco	Arista
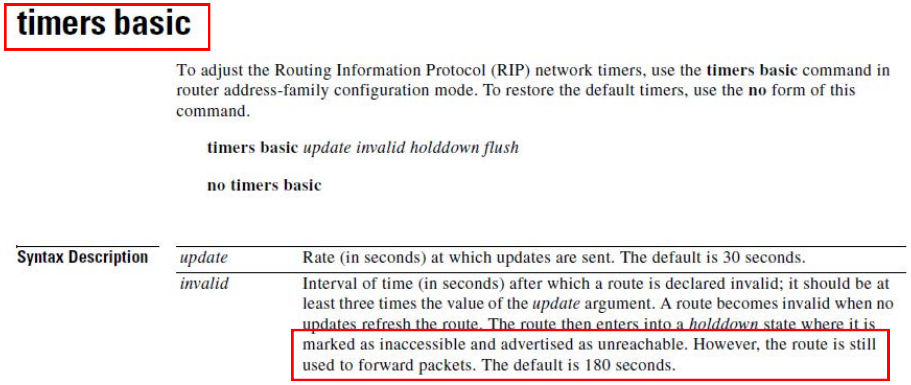 **timers basic** To adjust the Routing Information Protocol (RIP) network timers, use the **timers basic** command in router address-family configuration mode. To restore the default timers, use the **no** form of this command. **timers basic** *update invalid holddown flush* **no timers basic** Syntax Description *update* — Rate (in seconds) at which updates are sent. The default is 30 seconds. *invalid* — Interval of time (in seconds) after which a route is declared invalid; it should be at least three times the value of the *update* argument. A route becomes invalid when no updates refresh the route. The route then enters into a *holddown* state where it is marked as inaccessible and advertised as unreachable. However, the route is still used to forward packets. The default is 180 seconds. Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at L3-538	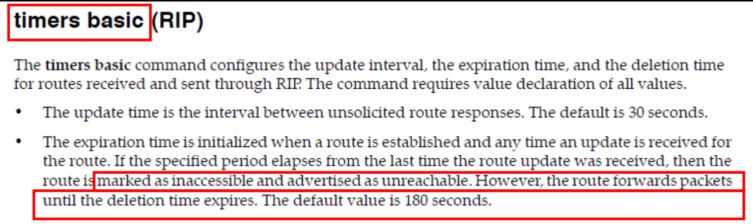 **timers basic (RIP)** The **timers basic** command configures the update interval, the expiration time, and the deletion time for routes received and sent through RIP. The command requires value declaration of all values. • The update time is the interval between unsolicited route responses. The default is 30 seconds. • The expiration time is initialized when a route is established and any time an update is received for the route. If the specified period elapses from the last time the route update was received, then the route is marked as inaccessible and advertised as unreachable. However, the route forwards packets until the deletion time expires. The default value is 180 seconds. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 1671
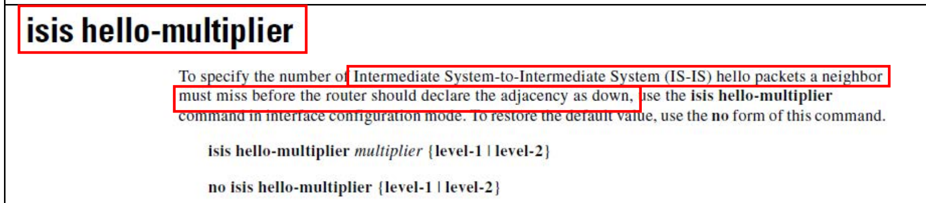 **isis hello-multiplier** To specify the number of Intermediate System-to-Intermediate System (IS-IS) hello packets a neighbor must miss before the router should declare the adjacency as down, use the **isis hello-multiplier** command in interface configuration mode. To restore the default value, use the **no** form of this command. **isis hello-multiplier** *multiplier* {**level-1** | **level-2**} **no isis hello-multiplier** {**level-1** | **level-2**} Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at L3-224	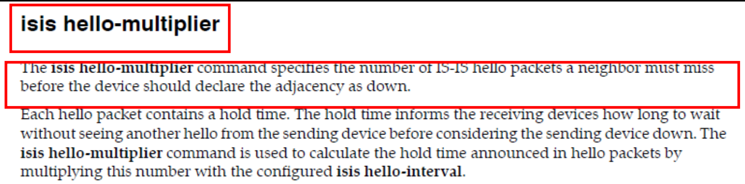 **isis hello-multiplier** The **isis hello-multiplier** command specifies the number of IS-IS hello packets a neighbor must miss before the device should declare the adjacency as down. Each hello packet contains a hold time. The hold time informs the receiving devices how long to wait without seeing another hello from the sending device before considering the sending device down. The **isis hello-multiplier** command is used to calculate the hold time announced in hello packets by multiplying this number with the configured **isis hello-interval**. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 1685
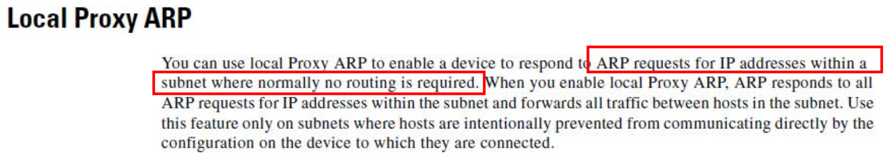 **Local Proxy ARP** You can use local Proxy ARP to enable a device to respond to ARP requests for IP addresses within a subnet where normally no routing is required. When you enable local Proxy ARP, ARP responds to all ARP requests for IP addresses within the subnet and forwards all traffic between hosts in the subnet. Use this feature only on subnets where hosts are intentionally prevented from communicating directly by the configuration on the device to which they are connected. Cisco NX-OS Unicast Routing Configuration Guide (2008), Release 4.0, at 2-5	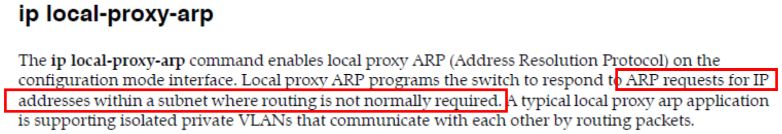 **ip local-proxy-arp** The **ip local-proxy-arp** command enables local proxy ARP (Address Resolution Protocol) on the configuration mode interface. Local proxy ARP programs the switch to respond to ARP requests for IP addresses within a subnet where routing is not normally required. A typical local proxy arp application is supporting isolated private VLANs that communicate with each other by routing packets. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 1276

Cisco	Arista
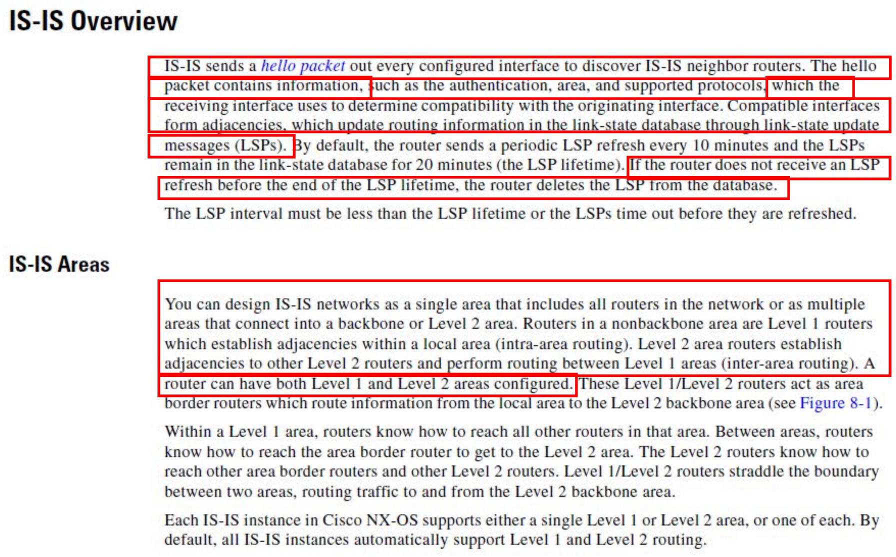 **IS-IS Overview** IS-IS sends a *hello packet* out every configured interface to discover IS-IS neighbor routers. The hello packet contains information, such as the authentication, area, and supported protocols, which the receiving interface uses to determine compatibility with the originating interface. Compatible interfaces form adjacencies, which update routing information in the link-state database through link-state update messages (LSPs). By default, the router sends a periodic LSP refresh every 10 minutes and the LSPs remain in the link-state database for 20 minutes (the LSP lifetime). If the router does not receive an LSP refresh before the end of the LSP lifetime, the router deletes the LSP from the database. The LSP interval must be less than the LSP lifetime or the LSPs time out before they are refreshed. **IS-IS Areas** You can design IS-IS networks as a single area that includes all routers in the network or as multiple areas that connect into a backbone or Level 2 area. Routers in a nonbackbone area are Level 1 routers which establish adjacencies within a local area (intra-area routing). Level 2 area routers establish adjacencies to other Level 2 routers and perform routing between Level 1 areas (inter-area routing). A router can have both Level 1 and Level 2 areas configured. These Level 1/Level 2 routers act as area border routers which route information from the local area to the Level 2 backbone area (see Figure 8-1). Within a Level 1 area, routers know how to reach all other routers in that area. Between areas, routers know how to reach the area border router to get to the Level 2 area. The Level 2 routers know how to reach other area border routers and other Level 2 routers. Level 1/Level 2 routers straddle the boundary between two areas, routing traffic to and from the Level 2 backbone area. Each IS-IS instance in Cisco NX-OS supports either a single Level 1 or Level 2 area, or one of each. By default, all IS-IS instances automatically support Level 1 and Level 2 routing. Cisco NX-OS Unicast Routing Configuration Guide, Release 4.0, at 8-2	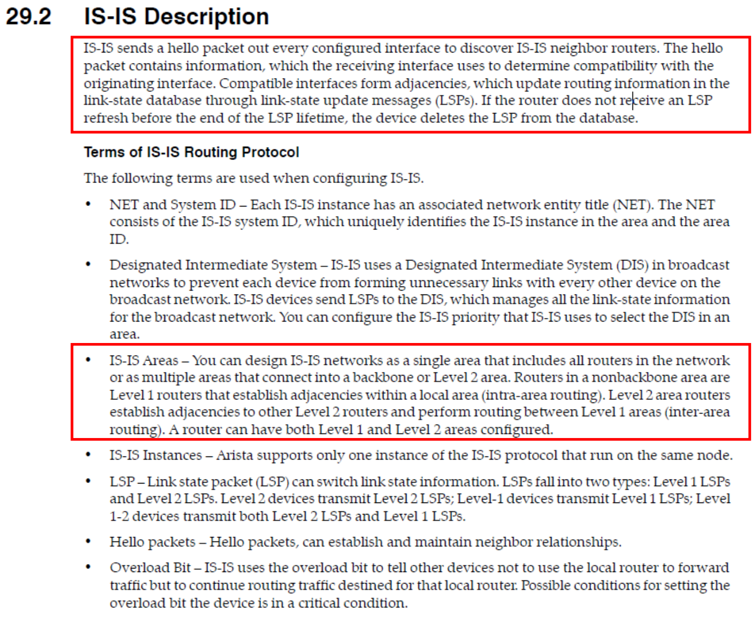 **29.2 IS-IS Description** IS-IS sends a hello packet out every configured interface to discover IS-IS neighbor routers. The hello packet contains information, which the receiving interface uses to determine compatibility with the originating interface. Compatible interfaces form adjacencies, which update routing information in the link-state database through link-state update messages (LSPs). If the router does not receive an LSP refresh before the end of the LSP lifetime, the device deletes the LSP from the database. **Terms of IS-IS Routing Protocol** The following terms are used when configuring IS-IS. • NET and System ID – Each IS-IS instance has an associated network entity title (NET). The NET consists of the IS-IS system ID, which uniquely identifies the IS-IS instance in the area and the area ID. • Designated Intermediate System – IS-IS uses a Designated Intermediate System (DIS) in broadcast networks to prevent each device from forming unnecessary links with every other device on the broadcast network. IS-IS devices send LSPs to the DIS, which manages all the link-state information for the broadcast network. You can configure the IS-IS priority that IS-IS uses to select the DIS in an area. • IS-IS Areas – You can design IS-IS networks as a single area that includes all routers in the network or as multiple areas that connect into a backbone or Level 2 area. Routers in a nonbackbone area are Level 1 routers that establish adjacencies within a local area (intra-area routing). Level 2 area routers establish adjacencies to other Level 2 routers and perform routing between Level 1 areas (inter-area routing). A router can have both Level 1 and Level 2 areas configured. • IS-IS Instances – Arista supports only one instance of the IS-IS protocol that run on the same node. • LSP – Link state packet (LSP) can switch link state information. LSPs fall into two types: Level 1 LSPs and Level 2 LSPs. Level 2 devices transmit Level 2 LSPs; Level-1 devices transmit Level 1 LSPs; Level 1-2 devices transmit both Level 2 LSPs and Level 1 LSPs. • Hello packets – Hello packets, can establish and maintain neighbor relationships. • Overload Bit – IS-IS uses the overload bit to tell other devices not to use the local router to forward traffic but to continue routing traffic destined for that local router. Possible conditions for setting the overload bit the device is in a critical condition. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 1674
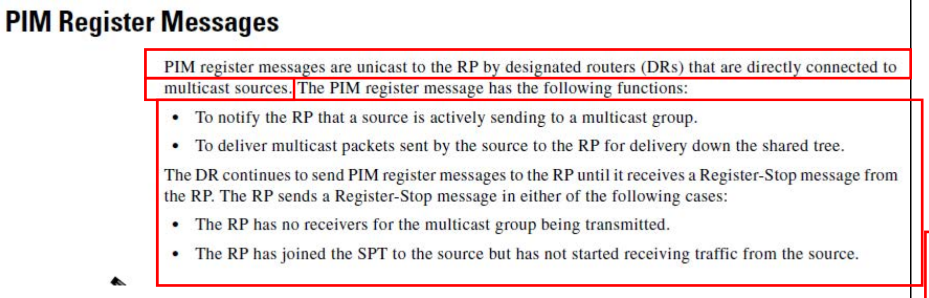 **PIM Register Messages** PIM register messages are unicast to the RP by designated routers (DRs) that are directly connected to multicast sources. The PIM register message has the following functions: • To notify the RP that a source is actively sending to a multicast group. • To deliver multicast packets sent by the source to the RP for delivery down the shared tree. The DR continues to send PIM register messages to the RP until it receives a Register-Stop message from the RP. The RP sends a Register-Stop message in either of the following cases: • The RP has no receivers for the multicast group being transmitted. • The RP has joined the SPT to the source but has not started receiving traffic from the source. Cisco NX-OS Multicast Routing Configuration Guide (2008), Release 4.0, at 3-7	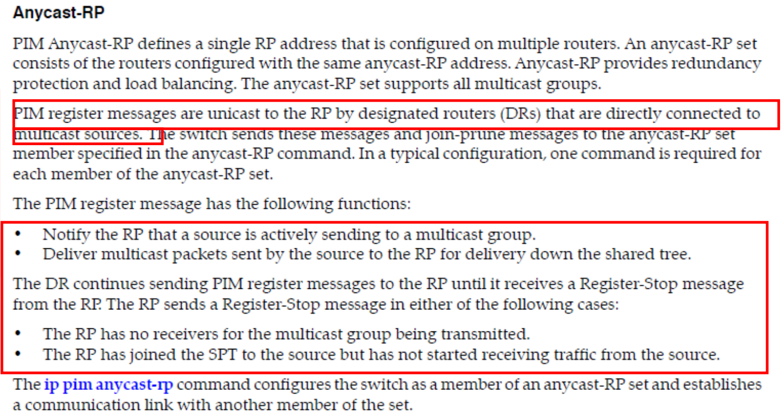 **Anycast-RP** PIM Anycast-RP defines a single RP address that is configured on multiple routers. An anycast-RP set consists of the routers configured with the same anycast-RP address. Anycast-RP provides redundancy protection and load balancing. The anycast-RP set supports all multicast groups. PIM register messages are unicast to the RP by designated routers (DRs) that are directly connected to multicast sources. The switch sends these messages and join-prune messages to the anycast-RP set member specified in the anycast-RP command. In a typical configuration, one command is required for each member of the anycast-RP set. The PIM register message has the following functions: • Notify the RP that a source is actively sending to a multicast group. • Deliver multicast packets sent by the source to the RP for delivery down the shared tree. The DR continues sending PIM register messages to the RP until it receives a Register-Stop message from the RP. The RP sends a Register-Stop message in either of the following cases: • The RP has no receivers for the multicast group being transmitted. • The RP has joined the SPT to the source but has not started receiving traffic from the source. The **ip pim anycast-rp** command configures the switch as a member of an anycast-RP set and establishes a communication link with another member of the set. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 1874

Cisco	Arista
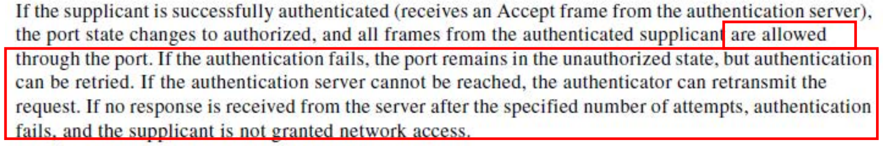 If the supplicant is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated supplicant are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the authenticator can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails, and the supplicant is not granted network access. Cisco DCNM Security Configuration Guide (2008), Release 4.0, at 6-5	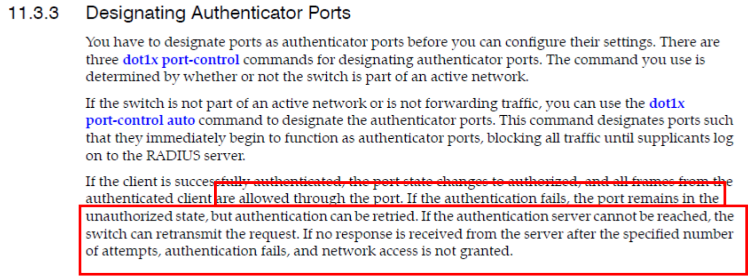 11.3.3 Designating Authenticator Ports You have to designate ports as authenticator ports before you can configure their settings. There are three **dot1x port-control** commands for designating authenticator ports. The command you use is determined by whether or not the switch is part of an active network. If the switch is not part of an active network or is not forwarding traffic, you can use the **dot1x port-control auto** command to designate the authenticator ports. This command designates ports such that they immediately begin to function as authenticator ports, blocking all traffic until supplicants log on to the RADIUS server. If the client is successfully authenticated, the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails, and network access is not granted. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 558
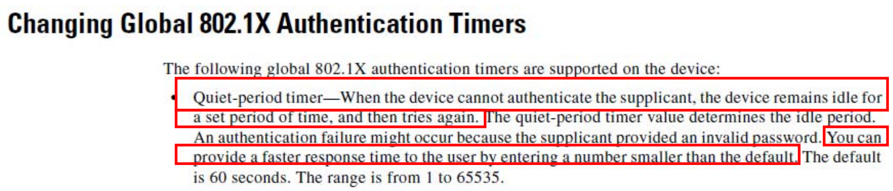 **Changing Global 802.1X Authentication Timers** The following global 802.1X authentication timers are supported on the device: - Quiet-period timer—When the device cannot authenticate the supplicant, the device remains idle for a set period of time, and then tries again. The quiet-period timer value determines the idle period. An authentication failure might occur because the supplicant provided an invalid password. You can provide a faster response time to the user by entering a number smaller than the default. The default is 60 seconds. The range is from 1 to 65535. Cisco DCNM Security Configuration Guide (2008), Release 4.0, at 6-14	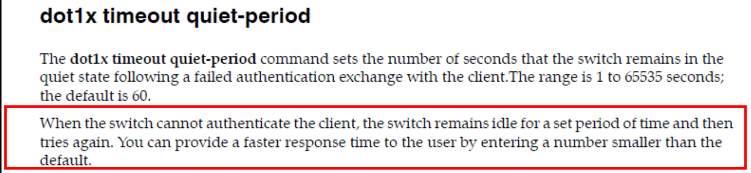 **dot1x timeout quiet-period** The **dot1x timeout quiet-period** command sets the number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client.The range is 1 to 65535 seconds; the default is 60. When the switch cannot authenticate the client, the switch remains idle for a set period of time and then tries again. You can provide a faster response time to the user by entering a number smaller than the default. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 569
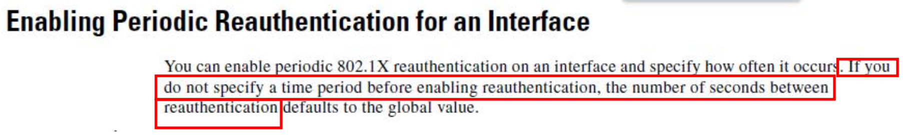 **Enabling Periodic Reauthentication for an Interface** You can enable periodic 802.1X reauthentication on an interface and specify how often it occurs. If you do not specify a time period before enabling reauthentication, the number of seconds between reauthentication defaults to the global value. Cisco DCNM Security Configuration Guide (2008), Release 4.0, at 6-14	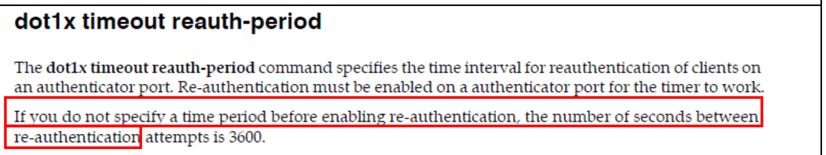 **dot1x timeout reauth-period** The **dot1x timeout reauth-period** command specifies the time interval for reauthentication of clients on an authenticator port. Re-authentication must be enabled on a authenticator port for the timer to work. If you do not specify a time period before enabling re-authentication, the number of seconds between re-authentication attempts is 3600. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 570

Cisco	Arista
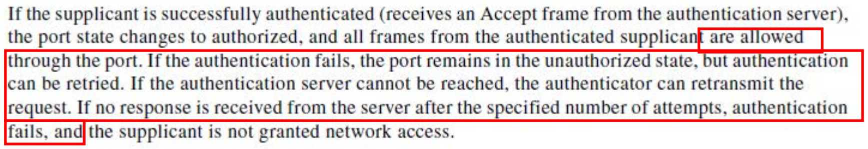 If the supplicant is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated supplicant are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the authenticator can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails, and the supplicant is not granted network access. Cisco NX-OS Security Configuration Guide (2008), Release 4.0, at 7-5	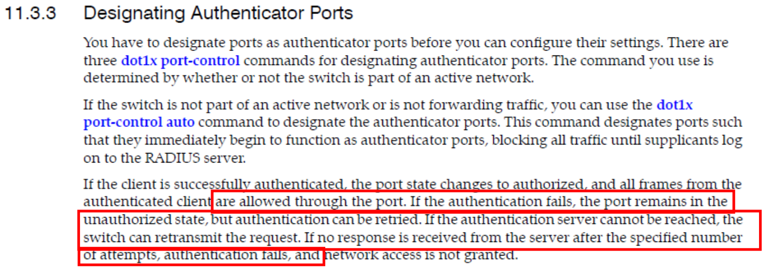 11.3.3 Designating Authenticator Ports You have to designate ports as authenticator ports before you can configure their settings. There are three **dot1x port-control** commands for designating authenticator ports. The command you use is determined by whether or not the switch is part of an active network. If the switch is not part of an active network or is not forwarding traffic, you can use the **dot1x port-control auto** command to designate the authenticator ports. This command designates ports such that they immediately begin to function as authenticator ports, blocking all traffic until supplicants log on to the RADIUS server. If the client is successfully authenticated, the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails, and network access is not granted. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 558
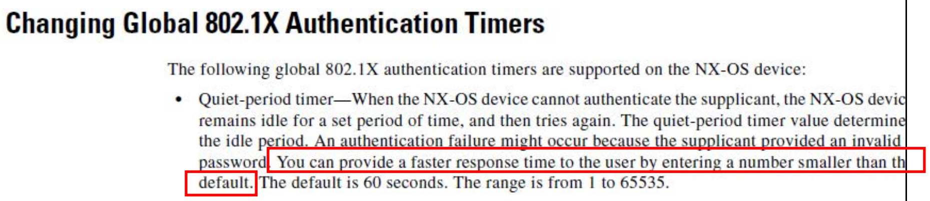 **Changing Global 802.1X Authentication Timers** The following global 802.1X authentication timers are supported on the NX-OS device: • Quiet-period timer—When the NX-OS device cannot authenticate the supplicant, the NX-OS devic remains idle for a set period of time, and then tries again. The quiet-period timer value determine the idle period. An authentication failure might occur because the supplicant provided an invalid password. You can provide a faster response time to the user by entering a number smaller than th default. The default is 60 seconds. The range is from 1 to 65535. Cisco NX-OS Security Configuration Guide (2008), Release 4.0, at 7-18	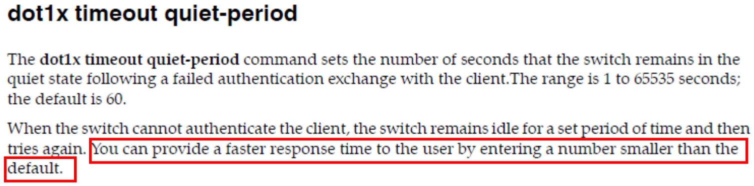 **dot1x timeout quiet-period** The **dot1x timeout quiet-period** command sets the number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client.The range is 1 to 65535 seconds; the default is 60. When the switch cannot authenticate the client, the switch remains idle for a set period of time and then tries again. You can provide a faster response time to the user by entering a number smaller than the default. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 569

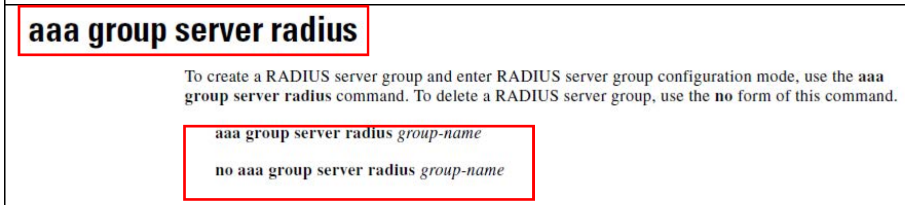
aaa group server radius
To create a RADIUS server group and enter RADIUS server group configuration mode, use the **aaa group server radius** command. To delete a RADIUS server group, use the **no** form of this command.
aaa group server radius *group-name*
no aaa group server radius *group-name*
 	
Cisco NX-OS Security Command Reference (2008), Release 4.0, at 17
	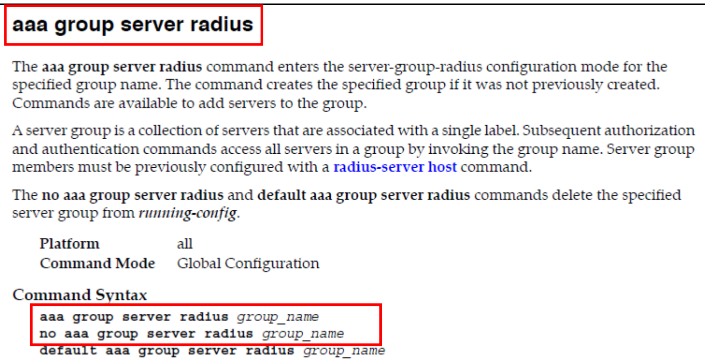
	aaa group server radius
	The **aaa group server radius** command enters the server-group-radius configuration mode for the specified group name. The command creates the specified group if it was not previously created. Commands are available to add servers to the group.
	A server group is a collection of servers that are associated with a single label. Subsequent authorization and authentication commands access all servers in a group by invoking the group name. Server group members must be previously configured with a **radius-server host** command.
	The **no aaa group server radius** and **default aaa group server radius** commands delete the specified server group from *running-config*.
	Platform all
	Command Mode Global Configuration
	Command Syntax
	`aaa group server radius group_name`
	`no aaa group server radius group_name`
	`default aaa group server radius group_name`
	
	Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 224

Cisco	Arista
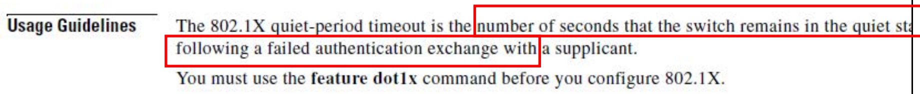 **Usage Guidelines** The 802.1X quiet-period timeout is the number of seconds that the switch remains in the quiet st following a failed authentication exchange with a supplicant. You must use the **feature dot1x** command before you configure 802.1X. Cisco NX-OS Security Command Reference (2008), Release 4.0, at 119	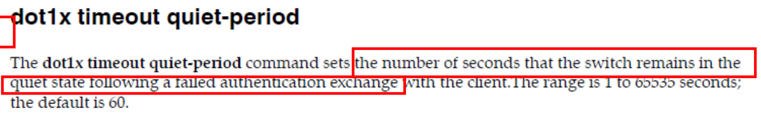 **dot1x timeout quiet-period** The **dot1x timeout quiet-period** command sets the number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client.The range is 1 to 65535 seconds; the default is 60. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 569
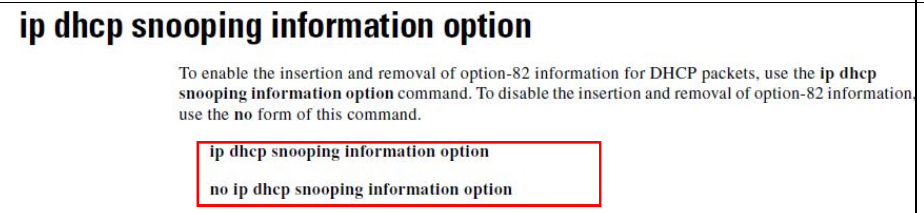 **ip dhcp snooping information option** To enable the insertion and removal of option-82 information for DHCP packets, use the **ip dhcp snooping information option** command. To disable the insertion and removal of option-82 information, use the **no** form of this command. **ip dhcp snooping information option** **no ip dhcp snooping information option** Cisco NX-OS Security Command Reference (2008), Release 4.0, at 196	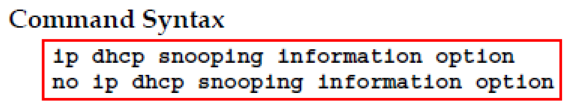 Command Syntax **`ip dhcp snooping information option`** **`no ip dhcp snooping information option`** Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 1270
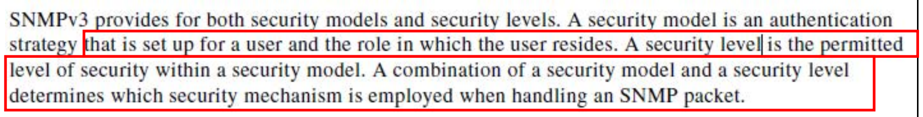 SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet. Cisco NX-OS System Management Configuration Guide (2008), Release 4.0, at 7-2	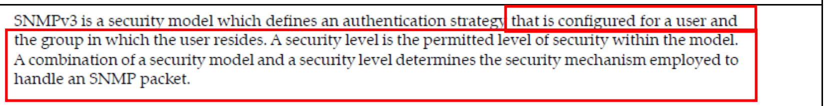 SNMPv3 is a security model which defines an authentication strategy that is configured for a user and the group in which the user resides. A security level is the permitted level of security within the model. A combination of a security model and a security level determines the security mechanism employed to handle an SNMP packet. Arista Configuration Guide v. 4.14.3F - Rev. 2 (10/2/14), at 1964